

Bayesian hierarchical modelling for inferring genetic interactions in yeast

Jonathan Heydari, Conor Lawless, David A. Lydall and Darren J. Wilkinson
Newcastle University, UK

Summary. Quantitative Fitness Analysis (QFA) is a high-throughput experimental and computational methodology for measuring the growth of microbial populations. QFA screens can be used to compare the health of cell populations with and without a mutation in a query gene in order to infer genetic interaction strengths genome-wide, examining thousands of separate genotypes. We introduce Bayesian, hierarchical models of population growth rates and genetic interactions that better reflect QFA experimental design than current approaches. Our new approach models population dynamics and genetic interaction simultaneously, thereby avoiding passing information between models via a univariate fitness summary. Matching experimental structure more closely, Bayesian hierarchical approaches use data more efficiently and find new evidence for genes which interact with yeast telomeres within a published dataset.

Keywords: Epistasis; Fitness; Genomic; Hierarchical; Interaction;

1. Introduction

There are many reasons to study the growth of microbes, including to prevent the growth of pathogenic bacteria or fungi and to encourage the growth of yeasts in industrial applications or during food production. Another reason is the study of eukaryotic microbes, such as the yeasts *Saccharomyces cerevisiae* and *Schizosaccharomyces pombe*, as biological models of cells in higher eukaryotes (e.g. of human cells).

Evolutionary fitness in a given environment: the probability of genetic material from an individual contributing to the gene pool of the next generation, is an important characteristic of a population that is optimised by natural selection. Rate of cell division is a major component of fitness, directly affecting the ability of individuals to compete for resources such as space and nutrients. By measuring and comparing the growth rates of microbial populations (cultures) we can assess and rank the fitness or health of such populations in a given environment or in a given genetic background.

Quantitative Fitness Analysis (QFA) is a method for measuring the growth and fitness of independent microbial cultures inoculated onto solid agar surfaces (Banks et al., 2012; Addinall et al., 2011). During QFA we inoculate cell cultures at densities of between 96 and 1536 cultures per plate of agar, repeatedly photographing cultures as they grow, converting photographs to quantitative estimates of cell density (Lawless et al., 2010). We summarise observations of increasing cell density with time (growth curves) by fitting population growth models to observed data. We use fitted model parameters, such as the intrinsic growth rate parameter of the logistic growth model, to define several measures of culture fitness (Addinall et al., 2011).

Quantifying the fitness of hundreds of strains on a single plate, under identical environmental conditions, allows a range of powerful experimental designs. Biological experiments examining the effect of a condition on selected populations, or the effect of selected conditions on one population, are often called screens. Screening independent replicate cultures with the same genotype allows us to measure biological heterogeneity and to capture technical error (which represents the effect of measurement error, fungal and bacterial contamination, positioning errors

and agar cracking in these experiments). Comparing cultures with different genotypes allows us to explore the relative importance of genes and gene products in a given environment or genetic background. An important reason for carrying out QFA is to compare the fitnesses of cultures with distinct genotypes in order to quantify the strength of interaction between genes (epistasis). Screening fitnesses and genetic interaction strengths on a genome-wide scale allows us to study the behaviour of gene products in living cells systematically. Ready-made genome-wide libraries of strains with distinct genotypes (each with an individual gene deleted, for example) are available and can be mated with selected strains to generate libraries targeted at particular biological processes of interest. A typical high-throughput, genome-wide QFA screen, examining the fitness of replicate cultures of 5,000 of different genotypes, includes hundreds of plates that are inoculated, photographed and incubated by laboratory robots.

The genome-wide QFA experiments that we re-analyse in this paper (see Section 4) were designed to inform us about telomere biology in eukaryotic cells. Telomeres are the ends of linear chromosomes found in most eukaryotic organisms (Greider and Blackburn, 1985), capping chromosome ends to ensure genetic stability, and are usually required for cells to progress through the cell cycle. Functional telomere caps help to prevent cancer and, since human telomeres shorten at each round of cell division (Olovnikov, 1973), some researchers claim that telomere-induced replicative senescence is an important component of human ageing. QFA experiments were carried out using *Saccharomyces cerevisiae* (brewer's yeast), a model eukaryotic organism widely used to study genetics. Yeasts are ideal for genome-wide analysis of gene function, as genetic modification of yeast cells is relatively straightforward and yeast cultures grow quickly; millions of yeast cells can be grown overnight, whereas the same number of human cells could take weeks to grow.

In these experiments, we used a genome-wide collection of *S. cerevisiae* strains, each carrying one of the set of about 5,000 single Open Reading Frame (ORF) deletions that are not essential for cell survival. An open reading frame is a DNA sequence containing no stop codons, which means that it has the potential to be translated into a protein or peptide. We refer to the mutations in this collection as *orf* Δ s; Δ is the standard genetics nomenclature for a deletion. Identifying ORFs from sequences is the first step in identifying genes, and using a library of ORFs allows the possibility of discovering biological function for sequences that were previously thought to be untranslated. However, the majority of ORFs in the collection we analyse have been confirmed as genes of known function and so *orf* Δ s are largely equivalent to gene deletions.

The strain collection was mated with a (query) background strain carrying the *cdc13-1* mutation, chosen for its relevance to telomere biology, to give a new library of strains carrying two mutations. Comparing fitnesses with a second, new library of strains, built from the deletion collection mated with a strain carrying a neutral control background mutation (*ura3* Δ) allows the separation of the effect of the *cdc13-1* mutation from that of deletions from the original collection.

More generally, we use QFA to infer genetic interaction strengths by comparing fitnesses in two QFA screens: a control screen and a query screen. All strains within a query screen differ from their control screen counterparts by a common condition such as a background gene mutation, drug treatment, temperature or other treatment. To identify strains that interact with the query condition we can compare the corresponding fitness responses for each strain in the library under the query and control conditions. Interactions with the query condition are identified by finding gene disruptions in the query screen whose fitnesses deviate significantly from those predicted by a theoretical model of genetic independence, given the fitness of corresponding gene disruptions in the control screen. Independent replicate cultures are inoculated and grown across several agar plates for each strain under each condition to capture biological heterogeneity and measurement error.

In the original analysis presented by Addinall et al. (2011), logistic models of population growth were fitted to observed cell density time courses by least squares, thereby generating a univariate fitness estimate for each time course. A linear model, predicting query strain fitness given control strain fitness, consistent with Fisher's multiplicative model of genetic independence, was used to test for genetic interaction between the query mutation and each deletion from the deletion collection. The significance of observed interactions was assigned using a simple frequentist linear modelling approach. A major limitation of the statistical model used in Addinall et al. (2011) is that it assumes that replicate culture fitness variances are the same for each $orf\Delta$. We expect that explicit modelling of heterogeneity will allow more robust identification of interactions, particularly where variability for a particular strain is unusually high (e.g. due to experimental difficulties).

Other large-scale quantitative genetic interaction screening approaches exist, such as E-MAP (Schuldiner et al., 2006) and SGA (Tong and Boone, 2006), but we expect QFA to provide higher quality fitness estimates by using a culture inoculation technique which results in a wider range of cell densities during culture growth and by capturing complete growth curves instead of using single time point assays (Lawless et al., 2010). QFA as presented by Addinall et al. (2011) and alternative genetic interaction screening approaches mentioned above use frequentist statistical methods that cannot account for all sources of experimental variation and do not partition variation into population, genotype and repeat levels. Further, the frequentist statistical approaches used in the methods above cannot incorporate prior beliefs.

With the Bayesian approach (Bernardo and Smith, 2007) that we adopt in this paper, we have more flexibility of model choice, allowing us to match model structure more closely to experimental design. Bayesian analysis allows us to use binary indicators to describe the evidence that each $orf\Delta$ interacts with the query mutation in terms of probability. Currently there is no standard frequentist approach which can deal with inference for a hierarchical model that simultaneously models logistic growth parameters and probability of genetic interaction. Using Bayesian hierarchical modelling (Zhang et al., 2014; Gelman and Hill, 2006) we look to extract as much information as possible from valuable QFA datasets.

Following the approach for determining epistasis from the comparison of two QFA screens presented by Addinall et al. (2011), we developed a two-stage approach to this problem: *i*) a hierarchical logistic growth curve model is fitted to cell density measurements to estimate fitness, then *ii*) fitness estimates are input to a hierarchical interaction model. Next, we developed a unified approach which we refer to as the joint hierarchical model (JHM). The JHM models mutant strain fitnesses and genetic interactions simultaneously, without having to pass information between two separate models. The JHM can also allow two important, distinct, microbial fitness phenotypes (population growth rate and carrying capacity) to provide evidence for genetic interaction simultaneously.

The paper is organized as follows: Section 2 describes the data from a typical QFA experiment. The two new models for Bayesian QFA are outlined in Section 3. In Section 4 the new Bayesian models are applied to a previously analysed QFA data set for identifying yeast genes interacting with a telomere defect. Section 5 discusses the relative merits of the newly developed Bayesian methods.

2. Defining Fitness

Observing changes in cell number in a microbial culture is the most direct way to estimate culture growth rate, an important component of microbial culture fitness. Direct counting of cells in a high-throughput experiment is not practical and so, during QFA, cell density estimates are made instead from culture photographs. Robotic assistance is required for both culture inoculation and

image capture during genome-wide screens which can include approximately 5,000 independent genotypes. We use estimates of the integrated optical density (IOD) generated by the image analysis tool Colonyzer (Lawless et al., 2010) to capture cell density dynamics in independent cultures during QFA (see Figure 1A).

Density estimates, scaled to normalise for camera resolution, are gathered for each culture and a dynamic model of population growth, the logistic model $\dot{x} = rx(1 - x/K)$ (Verhulst, 1845), is fit to the data. The logistic model ODE has three parameters: K , P and r , the carrying capacity (maximum achievable population density), culture inoculum density (initial condition) and culture growth rate respectively, and has the following analytical solution:

$$x(t; \theta) = \frac{KPe^{rt}}{K + P(e^{rt} - 1)}, \text{ where } P = x(0) \text{ and } \theta = (K, r, P). \quad (1)$$

This model describes self-limiting populations undergoing approximately exponential growth which slows as population density increases. During QFA, self-limited growth occurs because nutrients found in the solid agar substrate are consumed by the growing cell population. Ultimately the population density saturates at the carrying capacity once available nutrients are exhausted (see Figure 1).

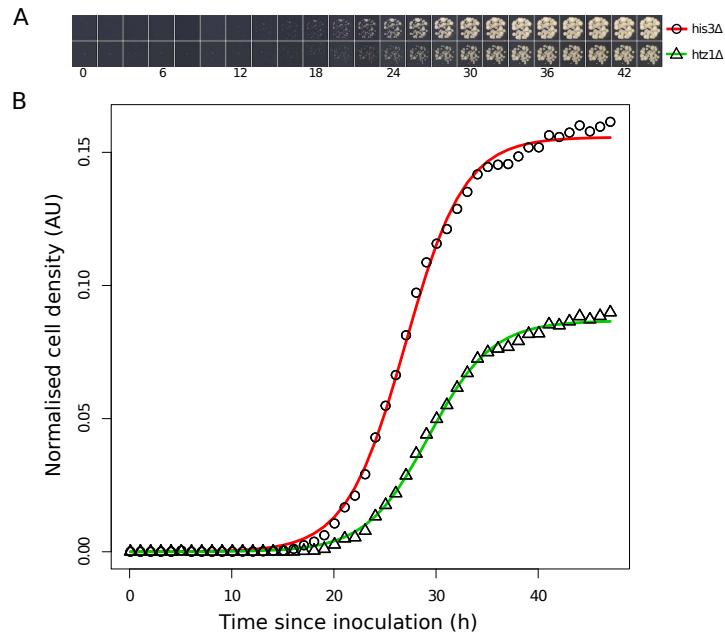


Fig. 1. QFA image data and growth curves. A) Timelapse images for two genetically modified *S. cerevisiae* cultures with different genotypes (indicated) corresponding to the time series measurements plotted in panel B. B) Timecourse cell density estimates derived from analysis of the timelapse images in panel A together with (least squares) fitted logistic growth curves.

We can construct several distinct, quantitative fitness measures based on fitted logistic model parameters. Addinall et al. (2011) present three univariate measures suitable for QFA: Maximum

Doubling Rate (D_R) and Maximum Doubling Potential (D_P), and their product $D_R \times D_P$,

$$\text{where } D_R = \frac{r}{\log\left(2\frac{K-P}{K-2P}\right)} \quad \text{and} \quad D_P = \frac{\log\left(\frac{K}{P}\right)}{\log(2)}. \quad (2)$$

D_R captures the rate at which microbes divide immediately after inoculation, when experiencing minimal intercellular competition or nutrient stress. A strain's growth rate largely dictates its ability to outcompete any neighbouring strains. D_P captures the number of divisions the culture is observed to undergo before saturation. A strain which can divide more often than its neighbours in a specific environment also has a competitive advantage.

The choice of a single overall fitness score depends on the aspects of microbial physiology most relevant to the biological question at hand. Typically the fitness definition $D_R \times D_P$ is used in QFA to account for both attributes simultaneously.

2.1. Epistasis

Epistasis is the phenomenon where the effects of one gene are modified by those of one or several other genes (Phillips, 1998). As presented in Addinall et al. (2011), here we use Fisher's multiplicative model of genetic independence (Cordell, 2002; Phenix et al., 2011) to represent the expected relationship between control strain fitness phenotypes and those of equivalent query strains in the absence of genetic interaction. We interpret genotypes for which the query strain fitness deviates significantly from this model of genetic independence as interacting significantly with the query mutation. Here, we use square bracket notation to represent a quantitative fitness measure. For example $[wt]$ and $[query]$ represent wild-type and query mutation fitnesses respectively. $orf\Delta$ is standard genetics nomenclature for the genotype of a strain with single gene (orf) deleted. We use this standard nomenclature to refer to an arbitrary strain from the deletion collection. We define new nomenclature to describe a strain containing two mutations. For example, $query : orf\Delta$ represents a strain with the query mutation along with an arbitrary single gene deletion. We use this nomenclature to refer to an arbitrary strain from the new query strain library constructed by crossing or mating a strain containing the query mutation with each of the strains in the genome-wide deletion collection. Fisher's multiplicative model of genetic independence can be written as follows:

$$[query : orf\Delta] \times [wt] = [query] \times [orf\Delta] \quad (3)$$

$$\Rightarrow [query : orf\Delta] = \frac{[query]}{[wt]} \times [orf\Delta]. \quad (4)$$

In (4), $\frac{[query]}{[wt]}$ is a constant for a given pair of QFA screens, meaning that if this model holds, there should be a linear dependence between $[query : orf\Delta]$ and $[orf\Delta]$ for all deletions $orf\Delta$. During genome-wide screens of thousands of independent $orf\Delta$ s we can assume that the majority of gene mutations in the library do not interact with the chosen query mutations. Therefore, even if the query or wild-type fitnesses are not available to us, we can still estimate the slope of this linear model by fitting it to all available fitness observations, before testing for strains which deviate significantly from the linear model. Any extra background condition, such as a gene mutation common to both the control and query strains (e.g. triple instead of double deletion strains for the query and control data sets), may change the biological interpretation of the interaction, but the same linear relationship is applicable. Besides the multiplicative model, there are other definitions for epistasis such as additive, minimum and log (Mani et al., 2008). Minimum is a suboptimal approach which may allow "masking" of interactions (Mani et al.,

2008). In this paper, we use a multiplicative interaction model (3), but we note that this is equivalent to an additive interaction model when looking at fitnesses on the log scale (Aylor and Zeng, 2008). Multiplicative and additive models are equivalent provided fitness data are scaled appropriately (Cordell, 2002).

2.2. Previous QFA methodology

Addinall et al. (2011) present QFA where the logistic growth model (1) is fit to experimental data by least squares to give parameter estimates (\hat{K}, \hat{r}) for each culture time course (each $orf\Delta$ replicate). Inoculum density is assumed known and the same across all $orf\Delta$ s and their repeats. After inoculating approximately 100 cells per culture, during the first several cell divisions there are so few cells that culture cell densities remain well below the detection threshold of cameras used for image capture and so, without sharing information across all $orf\Delta$ repeats, P cannot be estimated directly. It is therefore necessary to fix P to the same value for both screens, using an average estimate of P from preliminary least squares logistic growth model fits. Fitting the model to each $orf\Delta$ repeat separately means there is no sharing of information within an $orf\Delta$ or between $orf\Delta$ s when determining \hat{K} and \hat{r} .

Quantitative fitness scores ($F_{cm} = D_{R,cm} \times D_{P,cm}$) for each culture were defined (see (2) for definitions of D_R and D_P). The index c identifies the condition for a given $orf\Delta$: $c = 0$ for the control strain and $c = 1$ for the query strain. m identifies an $orf\Delta$ replicate. Scaled fitness measures \tilde{F}_{cm} are calculated for both the control and query screen such that the mean across all $orf\Delta$ s for a given screen is equal to 1. After scaling, any evidence that \tilde{F}_{0m} and \tilde{F}_{1m} are significantly different will be evidence of genetic interaction.

The linear model

$$\begin{aligned} \tilde{F}_{cm} &= \mu + \gamma_c + \varepsilon_{cm}, \text{ where } \gamma_0 = 0 \\ \varepsilon_{cm} &\sim N(0, \sigma^2), \text{ where } \varepsilon_{cm} \text{ is i.i.d.} \end{aligned} \quad (5)$$

was fitted to the control and query strain scaled fitness measure pairs for all unique $orf\Delta$ s in the gene deletion library. In (5), γ_1 represents the estimated strength of genetic interaction between the control and query strain. If the scaled fitnesses for the control and query strain are equivalent for a particular $orf\Delta$ such that they are both estimated by some μ , i.e. no evidence of genetic interaction, we would expect $\gamma_c = 0$. The model was fit by maximum likelihood, using the R function “*lmList*” (Pinheiro and Bates, 2000) with variation assumed to be the same for all strains in a given screen and the same for both control and query screens. Hence, for every gene deletion from the library an estimate of γ_1 was generated together with a p-value for whether it was significantly different from zero.

False discovery rate (FDR) corrected q-values were then calculated to determine levels of significance for each $orf\Delta$. Addinall et al. (2011) use the Benjamini-Hochberg test (Benjamini and Hochberg, 1995) for FDR correction. This test is commonly used in genomic analyses as although it assumes independence of test statistics, even if positive correlation exists between tests; the result is that FDR estimates are slightly conservative. Finally a list of $orf\Delta$ names, ranked by q-values, was output and $orf\Delta$ s with q-values below a significance cut-off of 0.05 were classed as showing significant levels of genetic interaction with the query mutation.

2.3. Random Effects Model

We attempted to improve on the Addinall et al. (2011) modelling approach within the frequentist paradigm by accounting for the hierarchical structure of the data with a random effects model (Pinheiro and Bates, 2000) of genetic interaction:

$$\begin{aligned}
f_{clm} &= \mu_c + Z_l + \gamma_{cl} + \varepsilon_{clm} \\
\mu_c &= \begin{cases} \mu + \alpha & \text{if } c = 0; \\ \mu & \text{if } c = 1, \end{cases} & \gamma_{cl} &= \begin{cases} 0 & \text{if } c = 0; \\ \gamma_l & \text{if } c = 1, \end{cases} \\
Z_l &\sim \mathcal{N}(0, \sigma_Z^2) & \varepsilon_{clm} &\sim \mathcal{N}(0, \sigma^2). \tag{6}
\end{aligned}$$

In the random effects model (REM) (6) and in models presented below, c identifies the condition for a given $orf\Delta$, l identifies a particular $orf\Delta$ from the gene deletion library and m identifies a repeat for a given $orf\Delta$. In (6) we use previously estimated F_{cm} to quantify interaction for all $orf\Delta$ s simultaneously. Introducing a random effect Z_l allows us to account for between subject variation by estimating a single σ_Z^2 . Unlike the Addinall et al. (2011) approach, we do not scale the observed values F_{clm} and instead introduce a parameter to model a condition effect μ_c . γ_{cl} represents the estimated strength of genetic interaction between an $orf\Delta$ and our query mutation. For a multiplicative model of epistasis we use an additive model to describe log transformed data $f_{clm} = \log(F_{clm} + 1)$, where F_{clm} are our observed fitnesses. We use the Benjamini-Hochberg test to correct for multiple testing in order to make a fair comparison with the Addinall et al. (2011) approach.

We find that $orf\Delta$ level variation in fitness cannot be modelled efficiently as random effects under the frequentist paradigm, which forces us to assume constant variance for all $orf\Delta$ s. The large number of random effects required (control and query observations for each of about 5,000 $orf\Delta$ s in a genome-wide screen) to model variances at the $orf\Delta$ level resulted in inference involving large matrix computations that either took too long to complete or were not possible using the computing hardware available to us. Similarly we found that it is not practical to model genetic interaction and cell population growth curves simultaneously as random effects under the frequentist paradigm. We attempted to model repeat level variation with a Normal distribution by fitting a model with a log-link function; however none of the non-linear model maximum likelihood algorithms we tried converged.

3. Bayesian hierarchical model inference

As an alternative to the maximum likelihood approach presented by Addinall et al. (2011) and the REM, we present a Bayesian, hierarchical methodology where *a priori* uncertainty about each parameter value is described by probability distributions (Bernardo and Smith, 2007) and information about parameter distributions is shared across $orf\Delta$ s and conditions. Plausible frequentist estimates from across 10 independent, unpublished QFA data sets, including a wide range of different background mutations and treatments were summarised to establish and quantify our *a priori* uncertainty in model parameters.

First and foremost, prior distributions describe our beliefs about parameter values. Priors should be at least diffuse enough to capture all plausible values (to capture the full range of observations in the datasets) and at least restrictive enough to rule out physically implausible values (to ensure efficient inference). Priors that are excessively vague are not consistent with the Bayesian paradigm and if they are unnecessarily diffuse can also result in computational difficulties during inference (see below for further details). The computational time required to overcome mixing problems from careless choice of prior distributions are likely to be considerable when fitting a large, hierarchical model to a rich dataset. Although using conjugate priors would allow slightly faster inference, we find that, for this particular application, the conjugate priors available for variance parameters (Gelman, 2006) are either too restrictive at low variance

(Inverse-gamma), not restrictive enough at low variance (half-t family of prior distributions) or are non-informative or largely discard the prior information available (Uniform). Here we have chosen the non-conjugate Log-normal distribution as a prior for precision parameters as we find that when appropriately parameterised the distribution reflects our prior beliefs about precision parameters and is only restrictive at extremely high and low variances.

We use three types of distribution to model parameter uncertainty: Log-normal, Normal and scaled t-distribution with three degrees of freedom. Particular care is needed in the choice of distributions for parameters which are in some sense close to the data, in order to ensure that the model is flexible enough to describe high-resolution datasets such as those captured during QFA. We use the Log-normal distribution to describe parameters which are required to be non-negative (e.g. parameters describing precisions, or repeat-level fitnesses) or parameter distributions which are found by visual inspection to be asymmetric. We use the Normal distribution to describe parameters which are symmetrically distributed (e.g. some prior distributions and the measurement error model) and we use the t-distribution to describe parameters whose uncertainty distribution is long-tailed (i.e. where using the Normal distribution would result in excessive shrinkage towards the mean). For example, after visual inspection of the variation of frequentist $orf\Delta$ level means about their population means in historical datasets, we found many unusually fit, dead or missing $orf\Delta$ s and concluded that $orf\Delta$ fitnesses would be well modelled by the t-distribution.

Instead of manually fixing the inoculum density parameter P as in Addinall et al. (2011) our Bayesian hierarchical models deal with the scarcity of information about the early part of culture growth curves by estimating a single P across all $orf\Delta$ s (and conditions in some of our models). Our new approach learns about P from the data and gives us a posterior distribution to describe our uncertainty about its value.

The new, hierarchical structure (Goldstein, 2011) implemented in our models reflects the structure of QFA experiments. Information is shared efficiently among groups of parameters, such as between repeat level parameters for a single mutant strain. Examples of the type of Bayesian hierarchical modelling which we use to model genetic interaction can be seen in Zhang et al. (2014) and Yi (2010), where hierarchical models are used to account for group effects.

In Phenix et al. (2011) the signal of genetic interaction is chosen to be “strictly ON or OFF” when modelling gene activity. We include this concept in our interaction models by using the posterior probability of a Bernoulli distributed indicator variable (O’Hara and Sillanpaa, 2009) to describe whether there is evidence of an $orf\Delta$ interacting with the query mutation; the more evidence of interaction, the closer posterior expectations will be to one.

Failing to account for all sources of variation within the experimental structure, such as the difference in variation between the control and query fitnesses may lead to inaccurate conclusions. By incorporating more information into the model with prior distributions and a more flexible modelling approach, we will increase statistical power. With an improved analysis it may then be possible for a similar number of genetic interactions to be identified with a smaller sample size (fewer replicate cultures), saving on the experimental costs associated with QFA.

Inference is carried out using Markov Chain Monte Carlo (MCMC) methods. The algorithm used is a Metropolis-within-Gibbs sampler where each full-conditional is sampled in turn either directly or using a simple Normal random walk Metropolis step. The scheme used is similar to that presented by Jow et al. (2014). Due to the large number of model parameters and large quantity of data from high-throughput QFA experiments, the algorithms used for carrying out inference often have poor mixing and give highly auto-correlated samples, requiring thinning. Posterior means are used to obtain point estimates where required.

In the following, we present a two-stage Bayesian, hierarchical modelling approach (Section 3.1 and 3.2) where we generate $orf\Delta$ fitness distributions and infer genetic interaction probabilities

separately. We then present a one-stage approach (Section 3.3) for inferring fitness and genetic interaction probabilities simultaneously. For the new approaches described in Section 3.1, 3.2 and 3.3 model fitting is carried out using the techniques discussed above, implemented in C for computational speed, and is freely available in the R package “qfaBayes” at <https://r-forge.r-project.org/projects/qfa>.

For the Bayesian models presented, the flow of information within the models and how each parameter is related to the data can be seen from the plate diagrams in Section 1 of the on-line supporting materials.

3.1. Separate Hierarchical Model

The separate hierarchical model (SHM), given in (7), models the growth of multiple yeast cultures using the logistic model described in (1), whose analytic solution is indicated by $x(t)$. The observational model at the time point level is given by

$$y_{lmn} \sim N(\hat{y}_{lmn}, (\nu_l)^{-1}) \quad \hat{y}_{lmn} = x(t_{lmn}; K_{lm}, r_{lm}, P),$$

where

$$\begin{aligned} l &= 1, 2, \dots, L && \text{orf}\Delta \text{ level} \\ m &= 1, \dots, M_l && \text{Repeat level} \\ n &= 1, 2, \dots, N_{lm} && \text{Time point level.} \end{aligned}$$

At the next level of the hierarchy (the repeat level), we have

$$\begin{aligned} \log K_{lm} &\sim N(K_l^o, (\tau_l^K)^{-1}) I_{(-\infty, 0]} & \log \tau_l^K &\sim N(\tau^{K,p}, (\sigma^{\tau,K})^{-1}) I_{[0, \infty)} \\ \log r_{lm} &\sim N(r_l^o, (\tau_l^r)^{-1}) I_{(-\infty, 3.5]} & \log \tau_l^r &\sim N(\tau^{r,p}, (\sigma^{\tau,r})^{-1}). \end{aligned}$$

Moving up, at the $\text{orf}\Delta$ level we have

$$\begin{aligned} e^{K_l^o} &\sim t(K^p, (\sigma^{K,o})^{-1}, 3) I_{[0, \infty)} & \log \sigma^{K,o} &\sim N(\eta^{K,o}, (\psi^{K,o})^{-1}) \\ e^{r_l^o} &\sim t(r^p, (\sigma^{r,o})^{-1}, 3) I_{[0, \infty)} & \log \sigma^{r,o} &\sim N(\eta^{r,o}, (\psi^{r,o})^{-1}) \\ \log \nu_l &\sim N(\nu^p, (\sigma^\nu)^{-1}) & \log \sigma^\nu &\sim N(\eta^\nu, (\psi^\nu)^{-1}). \end{aligned}$$

Finally, at the population level, we take

$$\begin{aligned} \log K^p &\sim N(K^\mu, (\eta^{K,p})^{-1}) & \log r^p &\sim N(r^\mu, (\eta^{r,p})^{-1}) \\ \log P &\sim N(P^\mu, (\eta^P)^{-1}) & \nu^p &\sim N(\nu^\mu, (\eta^{\nu,p})^{-1}) \\ \tau^{K,p} &\sim N(\tau^{K,\mu}, (\eta^{\tau,K,p})^{-1}) & \log \sigma^{\tau,K} &\sim N(\eta^{\tau,K}, (\psi^{\tau,K})^{-1}) \\ \tau^{r,p} &\sim N(\tau^{r,\mu}, (\eta^{\tau,r,p})^{-1}) & \log \sigma^{\tau,r} &\sim N(\eta^{\tau,r}, (\psi^{\tau,r})^{-1}). \end{aligned} \quad (7)$$

Dependent variable observations y_{lmn} (scaled cell density measurements) and independent variable t_{lmn} (time since inoculation) are model inputs, where n indicates the time point for a given $\text{orf}\Delta$ repeat. A directed, acyclic graph (DAG) for this model can be seen in Section 1 of the supporting on-line information. In this first hierarchical model, the logistic model is fit to query and control data separately.

In order to measure the variation between $\text{orf}\Delta$ s, parameters (K^p, σ_o^K) and (r^p, σ_o^r) are included at the population level of the hierarchy. Within- $\text{orf}\Delta$ variation is modelled by each set of

orf Δ level parameters (K_l^o, τ_l^K) and (r_l^o, τ_l^r) . Learning about these higher level parameters allows information to be shared across parameters lower in the hierarchy. A three-level hierarchical model is applied to (K, K_l^o, K_{lm}) and (r, r_l^o, r_{lm}) , sharing information on the repeat level and the *orf* Δ level. Note that *orf* Δ level parameters K_l^o and r_l^o are on the log scale ($e^{K_l^o}$ and $e^{r_l^o}$ are on the scale of the observed data).

Assuming a Normal error structure, random measurement error is modelled by the ν_l parameters (one for each *orf* Δ). Information on random error is shared across all *orf* Δ s by drawing $\log \nu_l$ from a Normal distribution parameterised by (ν_p, σ^{ν}) . A two-level hierarchical structure is also used for both the τ_l^K and τ_l^r parameters.

Modelling logistic model parameter distributions on the log scale ensures that parameters values remain strictly positive (a realistic biological constraint). Truncating distributions allows us to implement further, realistic constraints on the data. Truncating $\log r_{lm}$ values greater than 3.5 corresponds to disallowing biologically unrealistic culture doubling times faster than about 30 minutes and truncating of repeat level parameters $\log K_{lm}$ above 0 ensures that no carrying capacity estimate is greater than the maximum observable cell density, which is 1 after scaling.

orf Δ level parameters $e^{K_l^o}$ and $e^{r_l^o}$ are on the same scale as the observed data. Realistic biological constraints (positive logistic model parameters) are enforced at the repeat level; however both $e^{K_l^o}$ and $e^{r_l^o}$, which are assumed to have scaled t distributions, are truncated below zero to keep exponentiated parameters strictly positive.

Identifiability problems can arise for parameters K_{lm} and r_{lm} when observed cell densities are low and unchanging (consistent with growth curves for cultures which are very sick, dead or missing). In these cases, either K_{lm} or r_{lm} can take values near zero, allowing the other parameter to take any value without significantly affecting the model fit. In the Addinall et al. (2011) approach identification problems are handled in an automated post-processing stage: for cultures with low K estimates (classified as dead), r is automatically set to zero. Computing time wasted on such identifiability problems is reduced by truncating repeat level parameters r_{lm} , preventing the MCMC algorithms from becoming stuck in extremely low probability regions when K_{lm} takes near zero values. Similarly, $\log \tau_l^K$ parameters are truncated below 0 to overcome identifiability problems between parameters K_{lm} and r_{lm} when r_{lm} takes near zero values.

The SHM (7) is fit to both the query and control strains separately. Means are taken to summarise logistic growth parameter posterior distributions. Summaries $(\hat{K}_{lm}, \hat{r}_{lm}, \hat{P})$ for each *orf* Δ repeat are converted to univariate fitnesses F_{clm} where c identifies the condition (query or control), with any given fitness measure e.g. $D_R \times D_P$ (see (2) and Addinall et al. (2011)).

3.2. Interaction Hierarchical Model

After the SHM fit, the interaction hierarchical model (IHM), given in (8) can then be used to model estimated fitness scores F_{clm} and determine, for each *orf* Δ , whether there is evidence for interaction.

$c = 0, 1$	Condition level
$l = 1, \dots, L_c$	<i>orf</i> Δ level
$m = 1, \dots, M_{cl}$	Repeat level
Repeat level	

$$F_{clm} \sim N(\hat{F}_{cl}, (\nu_{cl})^{-1}) \quad \hat{F}_{cl} = e^{\alpha_c + Z_l + \delta_l \gamma_{cl}}$$

*orf*Δ level

$$\begin{aligned}
 e^{Z_l} &\sim t(Z^p, (\sigma^Z)^{-1}, 3) I_{[0, \infty)} & \log \sigma^Z &\sim N(\eta^Z, \psi^Z) \\
 \log \nu_{cl} &\sim N(\nu^p, (\sigma^\nu)^{-1}) & \log \sigma^\nu &\sim N(\eta^\nu, \psi^\nu) \\
 \delta_l &\sim Bern(p) \\
 e^{\gamma_{cl}} &= \begin{cases} 1 & \text{if } c = 0; \\ t(1, (\sigma^\gamma)^{-1}, 3) I_{[0, \infty)} & \text{if } c = 1. \end{cases} & \log \sigma^\gamma &\sim N(\eta^\gamma, (\psi^\gamma)^{-1})
 \end{aligned}$$

Condition level

$$\alpha_c = \begin{cases} 0 & \text{if } c = 0; \\ N(\alpha^\mu, \eta^\alpha) & \text{if } c = 1. \end{cases}$$

Population level

$$\log Z^p \sim N(Z^\mu, (\eta^{Z,p})^{-1}) \quad \nu^p \sim N(\nu^\mu, (\eta^{\nu,p})^{-1}) \quad (8)$$

F_{clm} are the observed fitness scores. A DAG for this model can be found in Section 1 of the supporting on-line materials. Fitnesses are passed to the IHM where query screen fitnesses are compared with control screen fitnesses, assuming genetic independence. Deviations from predicted fitnesses are evidence for genetic interaction. The interaction model accounts for between *orf*Δ variation with the set of parameters (Z^p, σ_Z) and within *orf*Δ variation by the set of parameters (Z_l, ν_l) . A linear relationship between the control and query *orf*Δ level parameters is specified with a scale parameter α_1 . Deviation from this relationship (genetic interaction) is accounted for by the term $\delta_l \gamma_{1l}$. A scaling parameter α_1 allows any effects due to differences in the control and query data sets to be scaled out, such as differences in genetic background, incubator temperature or inoculum density. The Bernoulli probability parameter p is our prior estimate for the probability of a given *orf*Δ showing evidence of genetic interaction. For the data set considered in Section 4 p is set to 0.05 as the experimenter's belief before the experiment was carried out was that 5% of the *orf*Δs would interact with the query. Observational noise is quantified by ν_{cl} . The ν_{cl} parameter accounts for difference in variation between condition i.e. the query and control data sets and for difference in variation between *orf*Δs.

The linear relationship between the control and query fitness scores, consistent with the multiplicative model of genetic independence, described in (4), is implemented in the IHM as $\hat{F} = e^{\alpha_c + Z_l + \delta_l \gamma_{cl}} = e^{\alpha_c} e^{Z_l + \delta_l \gamma_{cl}}$. Strains whose fitnesses lie along the linear relationship defined by the scalar α_1 show no evidence for interaction with the query condition. On the other hand, deviation from the linear relationship, represented by the posterior mean of $\delta_l \gamma_{1l}$ is evidence for genetic interaction. The larger the posterior mean for δ_l is, the higher the probability or evidence there is for interaction, while γ_{1l} is a measure of the strength of interaction. Where the query condition has a negative effect (i.e. decreases fitness on average, compared to the control condition), query fitnesses which are above and below the linear relationship are suppressors and enhancers of the fitness defect associated with the query condition respectively. A list of genes ranked by strength and direction of interaction with the query condition are ordered by the posterior means of $\delta_l \gamma_{cl}$. The *orf*Δs with $\hat{\delta}_l > 0.5$ are classified and labelled as showing "significant" evidence of interaction.

3.3. Joint Hierarchical Model

The joint hierarchical model (JHM), given in (9) is an alternative, fully Bayesian version of the two-stage approach described in Section 3.1 and 3.2.

$c = 0, 1$	Condition level
$l = 1, \dots, L_c$	$orf\Delta$ level
$m = 1, \dots, M_{cl}$	Repeat level
$n = 1, \dots, N_{clm}$	Time point level
Time point level	
$y_{clmn} \sim N(\hat{y}_{clmn}, (\nu_{cl})^{-1})$	$\hat{y}_{clmn} = x(t_{clmn}; K_{clm}, r_{clm}, P)$
Repeat level	
$\log K_{clm} \sim N(\alpha_c + K_l^o + \delta_l \gamma_{cl}, (\tau_{cl}^K)^{-1}) I_{(-\infty, 0]}$	$\log \tau_{cl}^K \sim N(\tau_c^{K,p}, (\sigma_c^{\tau,K})^{-1}) I_{[0, \infty)}$
$\log r_{clm} \sim N(\beta_c + r_l^o + \delta_l \omega_{cl}, (\tau_{cl}^r)^{-1}) I_{(-\infty, 3.5]}$	$\log \tau_{cl}^r \sim N(\tau_c^{r,p}, (\sigma_c^{\tau,r})^{-1})$
$orf\Delta$ level	
$e^{K_l^o} \sim t(K^p, (\sigma^{K,o})^{-1}, 3) I_{[0, \infty)}$	$\log \sigma^{K,o} \sim N(\eta^{K,o}, (\psi^{K,o})^{-1})$
$e^{r_l^o} \sim t(r^p, (\sigma^{r,o})^{-1}, 3) I_{[0, \infty)}$	$\log \sigma^{r,o} \sim N(\eta^{r,o}, (\psi^{r,o})^{-1})$
$\log \nu_{cl} \sim N(\nu^p, (\sigma^\nu)^{-1})$	$\log \sigma^\nu \sim N(\eta^\nu, (\psi^\nu)^{-1})$
$\delta_l \sim Bern(p)$	
$e^{\gamma_{cl}} = \begin{cases} 1 & \text{if } c = 0; \\ t(1, (\sigma^\gamma)^{-1}, 3) I_{[0, \infty)} & \text{if } c = 1. \end{cases}$	$\log \sigma^\gamma \sim N(\eta^\gamma, \psi^\gamma)$
$e^{\omega_{cl}} = \begin{cases} 1 & \text{if } c = 0; \\ t(1, (\sigma^\omega)^{-1}, 3) I_{[0, \infty)} & \text{if } c = 1. \end{cases}$	$\log \sigma^\omega \sim N(\eta^\omega, \psi^\omega)$
Condition level	
$\alpha_c = \begin{cases} 0 & \text{if } c = 0; \\ N(\alpha^\mu, \eta^\alpha) & \text{if } c = 1. \end{cases}$	$\beta_c = \begin{cases} 0 & \text{if } c = 0; \\ N(\beta^\mu, \eta^\beta) & \text{if } c = 1. \end{cases}$
$\tau_c^{K,p} \sim N(\tau^{K,\mu}, (\eta^{\tau,K,p})^{-1})$	$\log \sigma_c^{\tau,K} \sim N(\eta^{\tau,K}, (\psi^{\tau,K})^{-1})$
$\tau_c^{r,p} \sim N(\tau^{r,\mu}, (\eta^{\tau,r,p})^{-1})$	$\log \sigma_c^{\tau,r} \sim N(\eta^{\tau,r}, (\psi^{\tau,r})^{-1})$
Population level	
$\log K^p \sim N(K^\mu, (\eta^{K,p})^{-1})$	$\log r^p \sim N(r^\mu, (\eta^{r,p})^{-1})$
$\nu^p \sim N(\nu^\mu, (\eta^{\nu,p})^{-1})$	$\log P \sim N(P^\mu, (\eta^P)^{-1})$

Here, the dependent variable y_{clmn} (scaled cell density measurements) and independent variable t_{clmn} (time since inoculation) are input to the JHM. The JHM incorporates the key modelling ideas from both the SHM and the IHM with the considerable advantage that we can learn

about logistic growth model, fitness and genetic interaction parameters simultaneously, thereby avoiding having to choose a fitness measure or point estimates for passing information between models. The JHM is an extension of the SHM with the presence or absence of genetic interaction being described by a Bernoulli indicator and an additional level of error to account for variation due to the query condition. Genetic interaction is modelled in terms of the two logistic growth parameters K and r simultaneously.

By fitting a single JHM, we need only calculate posterior means, check model diagnostics and thin posterior samples once. However, the CPU time taken to reach convergence for any given data set is roughly twice that of the two-stage approach for a genome-wide QFA.

All of the SHM and IHM modelling assumptions described in Sections 3.1 and 3.2, such as distributional choices and hierarchical structure are inherited by the JHM. Similar to the interaction model in Section 3.2, linear relationships between control and query carrying capacity and growth rate (instead of fitness score) are assumed: $(e^{\alpha_c+K_l^o+\delta_l\gamma_{cl}}, e^{\beta_c+r_l^o+\delta_l\omega_{cl}})$.

4. Re-analysis of QFA experiments designed to learn about telomere biology

In this section we present a re-analysis of a previously published experiment, designed to inform us about the ways that eukaryotic cells respond to the loss of telomere caps that normally protect the ends of chromosomes from being erroneously recognised as a type of DNA damage. A pair of genome-wide QFA screens, were carried out in the model eukaryotic organism *S. cerevisiae* (brewer's yeast), comparing the fitness of control *ura3Δ* strains with query *cdc13-1* strains. These comparisons were made to identify genes that show evidence of interaction with the query mutation *cdc13-1*. Cdc13 is a *S. cerevisiae* protein which binds to telomeres and regulates telomere capping. *cdc13-1* is a temperature-sensitive allele of the *CDC13* gene. The ability of the altered Cdc13 protein produced by strains carrying the *cdc13-1* gene to cap telomeres is reduced at temperatures above 26 °C (Nugent et al., 1996), inducing a fitness defect that can be measured by QFA. The original experimental data used are freely available at <http://research.ncl.ac.uk/colonyzer/AddinallQFA/>. Addinall et al. (2011) present a list of inferred interaction strengths and p-values for significance of interaction, together with a fitness plot for this experiment.

Here, we will compare lists of genes classified as interacting with *cdc13-1* by the non-hierarchical frequentist approach presented by Addinall et al. (2011) and the hierarchical REM with those classified as interacting by our hierarchical Bayesian approaches.

4,294 non-essential *orfΔs* were selected from the yeast deletion collection and used to build the corresponding double deletion control and query strains. Independent replicate culture growth curves (time course observations of cell density) were captured for each control and query strain. The median and range for the number of replicates per *orfΔ* is 8 and [8, 144] respectively. The range for the number of time points for growth curves captured in the control experiment is [7, 22] and [9, 15] in the query experiment.

As in the Addinall et al. (2011) analysis, a list of 159 genes are stripped from our final list of genes for biological and experimental reasons. Priors for the models used throughout Section 4 are provided in Table 1. We have ensured that these priors are sufficiently diffuse to describe any QFA data set by inspecting 10 historical QFA data sets.

4.1. Model application

The Heidelberg-Welch (Heidelberger and Welch, 1981) and Raftery-Lewis (Raftery and Lewis, 1995) convergence diagnostics are used to determine whether convergence has been reached for all parameters. Posterior and prior densities are compared by eye to ensure that sample posterior

Table 1. Hyperparameter values specifying priors for the SHM, IHM and JHM.

SHM & JHM		SHM & JHM		JHM		IHM	
Parameter Name	Value	Parameter Name	Value	Parameter Name	Value	Parameter Name	Value
$\tau^{K,\mu}$	2.20	$\eta^{r,p}$	0.13	α^μ	0.00	Z_μ	3.66
$\eta^{\tau,K,p}$	0.02	ν^μ	19.82	η^α	0.25	$\eta^{Z,p}$	0.70
$\eta^{K,o}$	-0.79	$\eta^{\nu,p}$	0.02	β^μ	0.00	η^Z	0.10
$\psi^{K,o}$	0.61	P^μ	-9.04	η^β	0.25	ψ^Z	0.42
$\tau^{r,\mu}$	3.65	η^P	0.47	p	0.05	η^ν	0.10
$\eta^{\tau,r,p}$	0.02			η^γ	-0.79	ψ^ν	2.45
$\eta^{r,o}$	0.47			ψ^γ	0.61	ν^μ	2.60
$\psi^{r,o}$	0.10			η^ω	0.47	$\eta^{\nu,p}$	0.05
η^ν	-0.83			ψ^ω	0.10	α^μ	0.00
ψ^ν	0.86			$\eta^{\tau,K}$	2.20	η^α	0.31
K^μ	-2.01			$\psi^{\tau,K}$	0.02	p	0.05
$\eta^{K,p}$	0.03			$\eta^{\tau,r}$	3.65	η^γ	0.10
r^μ	0.97			$\psi^{\tau,r}$	0.02	ψ^γ	0.42

distributions are not restricted by the choice of prior distribution. ACF (auto-correlation) plot diagnostics are checked visually to ensure that serial correlation between sample values of the posterior distribution is low, ensuring that the effective sample size is similar to the actual sample size.

To assess how well the logistic growth model describes cell density observations we generate plots of raw data with fitted curves overlaid. Figure 2A, 2B and 2C show time series data for three different mutant strain repeats at 27°C, together with fitted logistic curves. Alternative fitness plots can be found in Section 3 of the online supporting material. We can see that each $orf\Delta$ curve fit represents the repeat level estimates well as each $orf\Delta$ level (red) curve lies in the region where most repeat level (black) curves are found. Sharing information between $orf\Delta$ s will also affect each $orf\Delta$ curve fit, increasing the probability of the $orf\Delta$ level parameters being closer to the population parameters. Comparing Figure 2A, 2B and 2C shows that the SHM captures heterogeneity at both the repeat and $orf\Delta$ levels.

Figure 2D demonstrates the hierarchy of information about the logistic model parameter K generated by the SHM for the *rad50* Δ control mutant strain (variation decreases going from population level down to repeat level). Figure 2D also shows that the posterior distribution for K is much more peaked than the prior, demonstrating that we have learned about the distribution of both the population and $orf\Delta$ parameters. Learning more about the repeat level parameters reduces the variance of our $orf\Delta$ level estimates. The posterior for the first time course repeat K_{clm} parameter shows exactly how much uncertainty there is for this particular repeat in terms of carrying capacity K .

4.1.1. Fitness Plots

Fitness plots are used to show which $orf\Delta$ s show evidence of genetic interaction. The plots are typically mean $orf\Delta$ fitnesses for query strains against the corresponding control strains.

Figure 3A is a fitness plot from Addinall et al. (2011) where growth curves and evidence for genetic interaction are modelled using the frequentist, non-hierarchical methodology discussed in Section 2.2. Figure 3B is a fitness plot for the frequentist hierarchical approach REM, described in 6, applied to the logistic growth parameter estimates used in Addinall et al. (2011). The number of genes identified as interacting with *cdc13-1* by Addinall et al. (2011) and by the REM are 715 and 315 respectively (Table 2). The REM has highlighted many strains which have low fitness. In order to fit a linear model to the fitness data and interpret results in terms of the multiplicative model we apply a log transformation to the fitnesses, thereby affecting the distribution of $orf\Delta$ level variation.

The REM accounts for between subject variation and allows for the estimation of a query

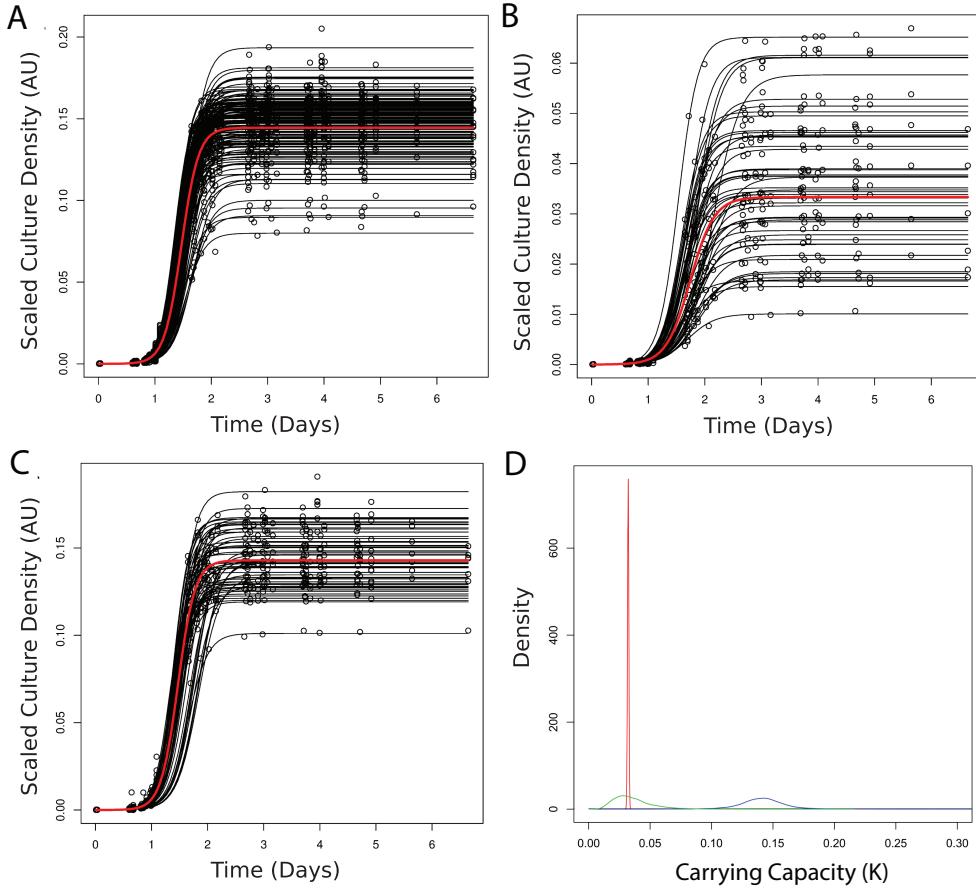


Fig. 2. Hierarchy of model fits and parameter estimates. Data for *orf* Δ repeats have been plotted in A, B and C, with SHM fitted curves overlaid in black for repeat level parameters and red for the *orf* Δ level parameter fit. A) SHM scatter plot for 144 *his3* Δ *ura3* Δ repeats at 27°C. B) SHM scatter plot for 48 *rad50* Δ *ura3* Δ repeats at 27°C. C) SHM scatter plot for 56 *exo1* Δ *ura3* Δ repeats at 27°C. D) SHM density plot of posterior predictive distributions for *rad50* Δ *ura3* Δ carrying capacity K hierarchy. The prior distribution for K^p is flat over this range. The posterior predictive for e^{K^p} is in blue and for K_{clm} in green. The posterior distribution of the first time course repeat K_{clm} parameter is in red. Parameters K^p , e^{K^p} and K_{clm} are on the same scale as the observed data.

mutation and *orf* Δ effect to be made simultaneously, unlike the model presented by Addinall et al. (2011). Due to the limitations of the frequentist hierarchical modelling framework, the REM model assumes equal variances for all *orf* Δ s and incorrectly describes *orf* Δ level variation as Log-normal, assumptions that are not necessary in our new Bayesian approaches.

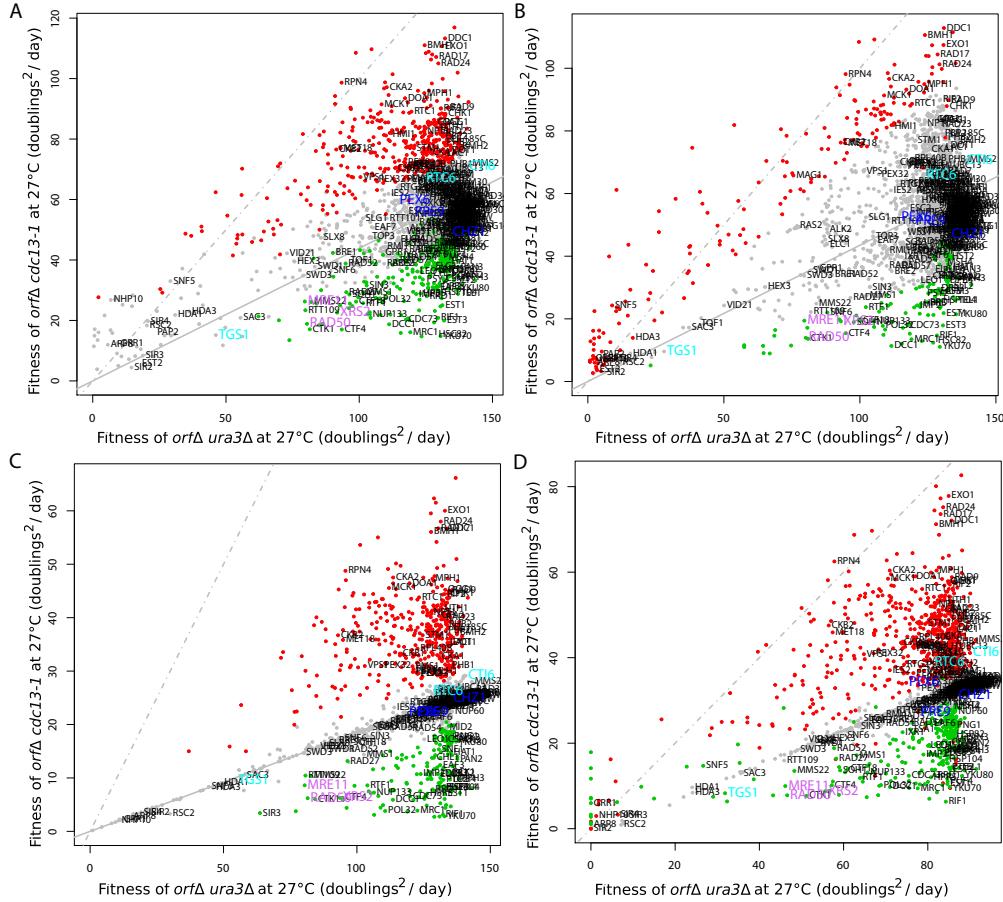


Fig. 3. Fitness plots comparing mean fitnesses ($F = D_R \times D_P$) for each $orf\Delta$ in a query and control screen. $orf\Delta$ s significantly suppressing or enhancing the *cdc13-1* fitness defect are highlighted in red and green respectively. A) Non-Bayesian, non-hierarchical fitness plot, based on Table S6 from Addinall et al. (2011). B) Non-Bayesian, hierarchical fitness plot, from fitting REM to data in Table S6 from Addinall et al. (2011). C) IHM fitness plot. D) JHM fitness plot, $orf\Delta$ s are classified as suppressors or enhancers based on analysis of growth parameter r : some strains are fitter in the query experiment than predicted based on control, but are classified as enhancers (green). A & B: significant interactors are classified as those with FDR corrected p-values < 0.05 . C & D: significant interactors have posterior probability $\delta_t > 0.5$. Labelled genes are annotated with GO terms from Table 2: “telomere maintenance”, “ageing”, “response to DNA damage stimulus” or “peroxisomal organization”, as well as genes identified as interactions using the JHM by considering K (see Figure 4) (blue) or by considering r (cyan) and the MRX complex genes (pink). Solid and dashed grey fitted lines are the line of equal fitness and linear model fits respectively.

Table 2. Number of genes interacting with *cdc13-1* at 27°C identified using each of four approaches: Add (Addinall et al., 2011), REM, IHM and JHM. Number of genes classified (or annotated) with four example GO terms (telomere maintenance, ageing, response to DNA damage stimulus and peroxisome organisation) are also listed. For the Add and REM approach, significant interactors are classified as those with FDR corrected p-values (q-values) < 0.05 . The label “half data” denotes analyses where only half of the available experimental observations are used.

Method	Suppressors			Enhancers			Hits			Suppressors (half data)			Enhancers (half data)			Hits (half data)			Telomere maintenance (N=33)			Aging (N=58)			Response to DNA damage stimulus (N=58)			Peroxisome organization (N=180)		
	Suppressors	Enhancers	Hits	Suppressors	Enhancers	Hits	Suppressors	Enhancers	Hits	Suppressors	Enhancers	Hits	Suppressors	Enhancers	Hits	Suppressors	Enhancers	Hits	Suppressors	Enhancers	Hits	Suppressors	Enhancers	Hits	Suppressors	Enhancers	Hits	Suppressors	Enhancers	Hits
Add	419	296	715	263	192	455	18	1.52E-06	0.0376	16	4.32E-05	0.1863	69	9.28E-12	8.14E-10	13	0.225	0.468												
REM	184	131	315	103	86	189	11	2.37E-05	0.0136	10	0.0004	0.0824	49	7.40E-16	1.73E-13	3	0.855	0.914												
IHM	404	172	576	252	113	365	14	6.57E-05	0.0051	16	0.0015	0.0445	55	4.60E-09	3.41E-07	10	0.318	0.524												
JHM	665	274	939	475	177	601	18	8.22E-05	0.0155	21	0.0015	0.0986	76	3.52E-09	1.99E-07	24	0.002	0.019												

4.2. Application of the two-stage modelling procedure to a suppressor/enhancer data set

Figure 3C is an IHM fitness plot with *orfΔ* level fitness measures generated using the new Bayesian two-stage methodology with fitness in terms of $D_R \times D_P$. 576 genes are identified by the IHM as genetic interactions (Table 2). Logistic parameter posterior means are used to generate fitness measures. For a gene (l) from the gene deletion library, (e^{Z_l}) is the fitness for the control and $(e^{\alpha_1 + Z_l + \delta_l \gamma_{c,l}})$ for the query. Similar to Figure 3A and 3B, Figure 3C shows how the majority of control strains are more fit than their query strain counterparts, with a mean fitted line lying below the line of equal fitness. Comparing the fitted lines in Figure 3A and 3B with Figure 3C, the IHM shows the largest deviation between the fitted line and the line of equal fitness, is largely due to the difference in P estimated with the SHM for the control and query data sets being scaled out by the parameter α_1 . If we fix P in our Bayesian models, as in the frequentist approach, genetic interactions identified are similar, but we then have the problem of choosing P . We recommend estimating P simultaneously with the other model parameters because if the choice of P is not close to the true value, growth rate r estimates must compensate and do not give accurate estimates for time courses with low carrying capacity K .

It can be seen that many of the interacting *orfΔs* have large deviations from the genetic independence line. This is because of the indicator variable in the model, used to describe genetic interaction. When there is enough evidence for interaction the binary variable is set to 1, otherwise it is set to 0. It is interesting to note that non-significant *orfΔs*, marked by grey points, lie amongst some of the significant strains. Many such points have high variance and we are therefore less confident that these interact with the query mutation. This feature of our new approach is an improvement over that presented in Addinall et al. (2011), which always shows evidence for an epistatic effect, for a given number of replicates, when mean distance from the genetic independence line is large, regardless of actual strain fitness variability.

4.3. Application of the Joint Hierarchical Model to a suppressor/enhancer data set

Figure 3D is a JHM $D_R \times D_P$ fitness plot using the new, unified Bayesian methodology. 939 genes are identified by the JHM as genetic interactions (Table 2). Posterior means of model parameters are used to obtain the following fitness measures. For a gene (l) from the gene deletion library, $(e^{K_l^o}, e^{r_l^o})$ are used to evaluate the fitness for the control and $(e^{\alpha_1 + K_l^o + \delta_l \gamma_{c,l}}, e^{\beta_1 + r_l^o + \delta_l \omega_{c,l}})$ for the

Table 3. Genes interacting with *cdc13-1* and GO terms over-represented in list of interactions according to each approach A) Number of genes identified for each approach (Add, REM, IHM and JHM) and the overlap between the approaches. 4135 genes from the *S. cerevisiae* single deletion library are considered. B) Number of GO terms identified for each approach and the overlap between the approaches. 6107 *S. cerevisiae* GO Terms available.

A.	REM:0				REM:1				B.	REM:0				REM:1			
	Add:0	Add:1	Add:0	Add:1	Add:0	Add:1	Add:0	Add:1		Add:0	Add:1	Add:0	Add:1	Add:0	Add:1	Add:0	Add:1
IHM:0	JHM:0	3097	54	31	10				IHM:0	JHM:0	5813	21	58	7			
	JHM:1	231	78	29	29					JHM:1	46	8	6	10			
IHM:1	JHM:0	1	2	1	0				IHM:1	JHM:0	20	15	3	12			
	JHM:1	30	327	0	215					JHM:1	13	54	2	147			

query.

Instead of producing a fitness plot in terms of $D_R \times D_P$, it can also be useful to analyse carrying capacity K and growth rate r fitness plots as, in the JHM, evidence for genetic interaction comes from both of these parameters simultaneously. Fitness plots in terms of logistic growth parameters are useful for identifying some unusual characteristics of $orf\Delta$ s. For example, an $orf\Delta$ may be defined as a suppressor in terms of K but an enhancer in terms of r . To enable direct comparison with the Addinall et al. (2011) analyses we generated a $D_R \times D_P$ fitness plot, Figure 3D.

4.4. Comparison with previous analysis

4.4.1. Significant genetic interactions

Of the genes identified as interacting with *cdc13-1* some are identified consistently across all four approaches (215 out of 1038, see Table 3A). Of the hits identified by the JHM (939), the majority (639) are common with those in the previously published Addinall et al. (2011) approach. However, 231 of 939 are uniquely identified by the JHM and could be interesting candidates for further study.

To examine the evidence for some interactions uniquely identified by the JHM in more detail we compared the growth curves for three examples from the group of interactions identified only by the JHM. These examples (*chz1* Δ , *pre9* Δ and *pex6* Δ) are genetic interactions which can be identified in terms of carrying capacity K , but not in terms of growth rate r (a unique feature of the JHM, see Figure 4). By observing the difference between the fitted growth curve (red) and the expected growth curve, given no interaction (green) in Figure 4A, 4B and 4C we test for genetic interaction. Since the expected growth curves in the absence of genetic interaction are not representative of either the data or the fitted curves on the repeat and $orf\Delta$ level, there is evidence for genetic interaction.

We chose a prior for the probability p of a gene interacting with the background mutation as 0.05, and explore the effect of alternate choices below. We therefore expected to find 215 genes interacting. Using the Bayesian models, for which a prior is applicable (IHM and JHM), we find more genes than expected (576 and 939 interactions respectively, Table 2), demonstrating that this dataset is sufficiently information rich to overcome prior expectations. The JHM identifies the highest proportion of genes as hits out of all methods considered, particularly identifying suppressors of *cdc13-1* (Table 2). In fact, the JHM identifies more hits than the Addinall et al. (2011) approach, even when constrained to using only half of the available data. An important advantage of our new Bayesian approaches is that we no longer have to choose a q-value threshold. For the Addinall et al. (2011) approach to have similar numbers of interactions to the JHM, a less stringent q-value threshold would have to be justified *a posteriori* by the experimenter.

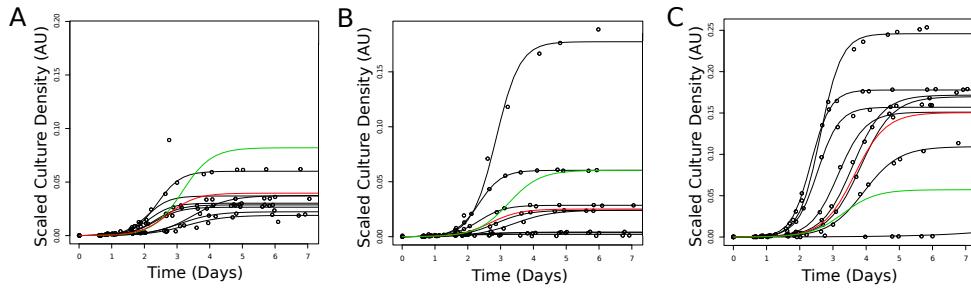


Fig. 4. Hierarchy of growth curve model fits for the JHM for some example genotypes. JHM data for *orfΔ* repeats have been plotted in A, B and C, with fitted curves overlaid in black for repeat level parameters, red for the *orfΔ* level query parameter fit and green for the expected *orfΔ* level query parameter fit with no genetic interaction. A) JHM scatter plot for 8 *chz1Δ cdc13-1* repeats. B) JHM scatter plot for 8 *pre9Δ cdc13-1* repeats. C) JHM scatter plot for 8 *pex6Δ cdc13-1* repeats.

4.4.2. Previously known genetic interactions

In order to compare the quality of our new, Bayesian hierarchical models with existing, frequentist alternatives, we examined the lists of genetic interactions identified by all the methods discussed and presented here. Comparing results with expected or previously known lists of interactions from the relevant literature, we find that genes coding for the MRX complex (*MRE11*, *XRS2* & *RAD50*), which are known to interact subtly with *cdc13-1* (Foster et al., 2006), are identified by all four approaches considered and can be seen in a similar position on all four fitness plots (Figure 3A, 3B, 3C and 3D).

By observing the genes labelled in Figure 3A and 3B we can see that the frequentist approaches are unable to identify many of the interesting genes identified by the JHM as these methods are unable to detect interactions for genes close to the genetic independence line. It seems likely that the JHM has extracted more information from deletion strain fitnesses observed with high variability than the Addinall et al. (2011) approach by sharing more information between levels of the hierarchy, consequently improving our ability to identify interactions for genes that are found closer to the line of genetic independence (subtle interactions). *CTI6*, *RTC6* and *TGS1* are three examples of subtle interactors identified only by the JHM (interaction in terms of *r* but not *K*) which all have previously known telomere-related functions (Franke et al., 2008; Keogh et al., 2005; Addinall et al., 2008).

We tested the biological relevance of results from the various approaches by carrying out unbiased Gene Ontology (GO) term enrichment analyses on the hits (lists of genes classified as having a significant interaction with *cdc13-1*) using the bioconductoR package GOstats (Falcon and Gentleman, 2007) (see Section 2 of the on-line supporting materials). As an example, fitness plots with genes co-annotated with the “telomere maintenance” highlighted can be seen in Section 3 of the on-line supporting materials.

Extracts from the list of top interactions identified by both the IHM and JHM are provided in Section 4 of the on-line supporting materials. Files including the full lists of genetic interactions for the IHM and JHM are also provided (<http://research.nci.ac.uk/qfa/HeydariQFABayes/>). Since we can use the JHM to identify interactions in terms of both *K* and *r* simultaneously, it is useful to order lists of suppressors and enhancers in terms of *K* and *r* as well as a fitness measure such as $D_R \times D_P$ for reviewing the results, see Section 5 of the on-line supporting materials.

All methods identify a large proportion of the genes in the yeast genome annotated with the

GO terms “telomere maintenance” and “response to DNA damage stimulus” (see Table 2 and the on-line supporting materials.), which were the targets of the original screen, demonstrating that they all correctly identify previously known hits of biological relevance. Interestingly, the JHM identifies many more genes annotated with the “ageing” GO term, which we also expect to be related to telomere biology (though the role of telomeres in ageing remains controversial) suggesting that the JHM is identifying novel, relevant interactions not previously identified by the Addinall et al. (2011) screen (see Table 2). Similarly, the JHM identifies a much larger proportion of the PEX “peroxisomal” complex (included in GO term: “peroxisome organisation”) as interacting with *cdc13-1* (see Table 2) including all of those identified in Addinall et al. (2011). Many of the PEX genes show large variation in both K and r , an example can be seen in Figure 4C for *pex6Δ*. Members of the PEX complex cluster tightly, above the fitted line in the fitness plot Figure 3D (fitness plots with highlighted genes for GO terms in Table 2 are given in Section 3 of the on-line supporting materials), demonstrating that although these functionally related genes are not strong interactors, the same behaviour is reproduced independently by multiple members of a known functional complex, suggesting that the predicted interactions are real. The results of tests for significant over-representation of all GO terms can be found online: <http://research.ncl.ac.uk/qfa/HeydariQFABayes/>.

Overall, within the lists of genes identified as interacting with *cdc13-1* by the Addinall et al. (2011), REM, IHM and JHM approaches, 274, 245, 266 and 286 GO terms were significantly over-represented respectively (out of 6235 possible GO terms, see Table 3B). 147 were common to all approaches and examples from the group of GO terms over-represented in the JHM analysis and not in the Addinall et al. (2011) analysis seem internally consistent (e.g. “peroxisome organisation” GO term) and consistent with the biological target of the screen, telomere biology (significant GO terms for genes identified only by the JHM are also included in the spreadsheet document provided in the on-line supporting materials).

A major advantage of the Bayesian approaches presented here over the Addinall et al. (2011) approach is the measure used for classifying significant interactions. Classifying interactions with a posterior estimate for δ_l (the probability that an interaction exists) greater than 0.5 as significant is less arbitrary than the traditional frequentist approach of classifying interactions with p-values less than 0.05 as significant. Examining how the number of over-represented GO terms found in lists of interactors varies with the classification threshold shows that the Bayesian JHM approach is also less sensitive to the precise threshold values used. Figure 5 shows that the number of over-expressed GO terms found among hits is relatively stable in the region of $\delta_l = 0.5$ for the JHM compared to the equivalent number in the region of $q = 0.05$ for the Addinall et al. (2011) approach. Significantly over-expressed GO terms were identified using the `hyperGTest` function in the `GOstats` R package. Note that the values used to classify whether a gene interacts with *cdc13-1* at 27C (q-value and δ respectively, red vertical lines as presented in Section 4.4) are not directly comparable; however the full range of possible cutoffs for both values are plotted. In particular, using the frequentist Addinall et al. (2011) approach, the number of over-expressed GO terms falls rapidly where $q < 0.05$. We tested whether this observation depended on our choice of the parameter p , which represents our prior expectation of the proportion of genes interacting with the query, by generating similar sensitivity plots for p between 0.01 and 0.2 (Section 6 of the on-line supporting materials). We observed similar profiles of over-expressed GO terms for all values of p tested.

Comparing the genetic interaction strengths generated by the Bayesian hierarchical models and frequentist analysis, we find that the results are largely similar (Section 6 of the on-line supporting materials); however the GO term analysis described above suggests that the differences are important.

The results of a simulation study comparing the sensitivity and specificity of the Addinall

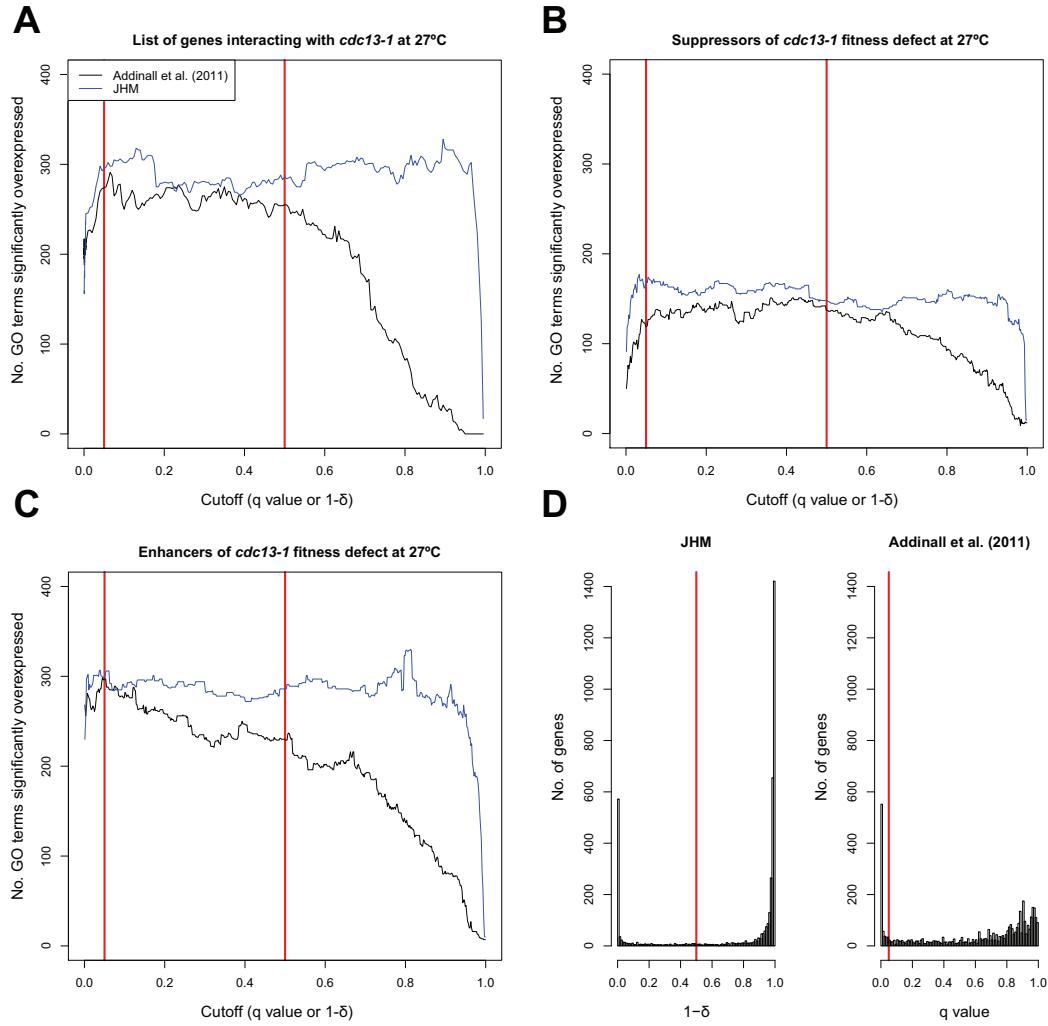


Fig. 5. Sensitivity to significance thresholds. A), B) & C) Comparison of the number of significantly over-expressed GO terms identified in lists of significant interactors found using the Addinall et al. (2011) method and using the JHM. D) Frequency histograms showing distributions of classifier values after looking for genes interacting with *cdc13-1* at 27°C using the Addinall et al. (2011) method or using the JHM.

et al. (2011) approach, the REM, the SHM and the JHM are summarised in section 8 of the supplementary materials. We find that the JHM correctly identified a higher proportion of “true” interactions in a synthetic dataset than the Addinall et al. (2011) approach.

4.4.3. Hierarchy and model parameters

The hierarchical structure and model choices included in the Bayesian JHM and IHM are derived from the known experimental structure of QFA. Different levels of variation for different

orf Δ s are expected and can be observed by comparing distributions of frequentist estimates or by visual inspection of yeast culture images. The direct relationship between experimental and model structure, together with the richness of detail and number of replicates included in QFA experimental design, reassures us that overfitting is not an issue in this analysis. For the *ura3* Δ 27°C and *cdc13-1* 27°C experiment with \sim 4,294 *orf* Δ s there are \sim 1.25 times the number of parameters in the JHM (\sim 200,000) compared to the two stage REM approach (\sim 160,000) but when compared to the large number of pairs of data points (\sim 830,000) there are sufficient degrees of freedom to justify our proposed Bayesian models.

4.4.4. Computing requirements

Our Bayesian hierarchical models require significant computational time. As expected, the mixing of chains in our models is weakest at population level parameters such as K_p and α_c . For the *ura3* Δ 27°C and *cdc13-1* 27°C dataset, running with an MCMC burn-in of 800,000 updates, followed by generating 1,000 samples thinned by a factor of 100, the JHM takes \sim 4 weeks to converge and produce a sufficiently large sample. The two stage Bayesian approach takes one week (with the IHM part taking \sim 1 day), whereas the REM takes \sim 3 days and the Addinall et al. (2011) approach takes \sim 3 hours. A QFA experiment can take over a month from start to finish and so analysis time is acceptable in comparison to the time taken for the creation of the data set but still a notable inconvenience. We expect that with further research effort, computational time can be decreased by using an improved inference scheme and that inference for the JHM could be completed in less than a week without parallelisation.

5. Discussion

We have joined a hierarchical model of microbial growth with a model for genetic interaction in order to learn about strain fitnesses, evidence for genetic interaction and interaction strengths simultaneously. By introducing Bayesian methodology to QFA we have been able to model the hierarchical nature of the experiment and expand the multiplicative model for genetic interaction to incorporate many sources of variation that previously had to be ignored.

We propose two new Bayesian hierarchical models to replace the current statistical analysis for identifying genetic interactions within a QFA screen comparison. The two-stage approach fits the SHM followed by the IHM, with univariate point estimate fitness definitions generated as an intermediate step. The two-stage approach can therefore be regarded as a Bayesian hierarchical version of the Addinall et al. (2011) approach. In contrast, the one-stage approach fits the JHM, which does not require a separate definition of fitness, allowing interaction to be identified by either growth rate (logistic parameter r) or final biomass achievable (logistic parameter K) by a given genotype. Our one-stage approach is a new method for detecting genetic interaction that further develops the interpretation of epistasis within QFA screens.

We present a hierarchical, frequentist approach using random effects, namely the REM, in an attempt to improve on the Addinall et al. (2011) approach. Due to the lack of flexibility in modelling assumptions allowable, the REM is unsuitable for modelling the distribution of *orf* Δ level variation or for simultaneously modelling genetic interaction and logistic growth curves.

The data from which logistic parameter estimates are derived during QFA are the result of a technically challenging, high-throughput experimental procedure with a diverse range of possible technical errors. Our Bayesian, hierarchical models allow us the flexibility to make distributional assumptions that more closely match the data. This allows us to switch between modelling parameter uncertainty with Normal, Log-Normal and Student's t distribution where appropriate.

QFA experimental design is intrinsically multilevel and is therefore more closely modelled by our hierarchical scheme. Consequently the JHM and IHM capture sources of variation not considered by Addinall et al. (2011). By sharing information across levels in the hierarchy, our models have allowed us to learn more about *orf* Δ s with weaker genetic interaction. Our more flexible model of variance also avoids misclassification of individual genotypes with high variance as having significant interactions. Without fully accounting for the variation described in the Bayesian hierarchical models, the previous Addinall et al. (2011) approach may have relatively poor power to detect subtle interactions, obscuring potential novel observations.

Many subtle, interesting genetic interactions may remain to be identified in the data from the QFA experiments we re-analyse in this paper. The JHM is better able to identify subtle interactions. For example, strains with little evidence for interaction with a background mutation in terms of growth rate but with strong evidence of interaction in terms of carrying capacity are sometimes classified as interactors using the JHM (see Figure 4). In our two-stage approaches, univariate fitness measures such as $D_R \times D_P$ are used in the intermediate steps, occasionally causing interaction in terms of one parameter to be masked by the other.

As expected, many genes previously unidentified by Addinall et al. (2011) have been identified as showing evidence of interaction using both of our Bayesian hierarchical modelling approaches. Genes which have been identified only by the JHM (see Figure 3D), such as those showing interaction only in terms of r , are found to be related to telomere biology in the literature. Currently there is not sufficient information available to identify the proportion of identified interactions that are true hits and so we use unbiased GO term enrichment analyses to confirm that the lists of genetic interactions closely reflect the true underlying biology. GO term annotations relevant to telomere biology are available for well-studied genes in the current literature. Unsurprisingly all of the approaches considered closely reflect the most well-known GO terms (see Table 2).

Computational time for the new Bayesian approach ranges from one to four weeks for one of the datasets presented in Addinall et al. (2011). This is of the same magnitude as the time taken to design and execute the experimental component of QFA (approximately six weeks).

Overall we recommend a JHM or “Bayesian QFA” for analysis of current and future QFA data sets as it accounts for more sources of variation than the Addinall et al. (2011) QFA methodology. With the JHM we have outlined new genes with significant evidence of interaction in the *ura3* Δ 27°C and *cdc13-1* 27°C experiment. The new Bayesian hierarchical models we present here will also be suitable for identifying new genes showing evidence of genetic interaction in backgrounds other than telomere activity. We hope that further, reductionist lab work by experimental biologists will give additional insight into the mechanisms by which the new genes we have uncovered interact with the telomere.

Acknowledgements

This research was supported by grants from the Biotechnology and Biological Sciences Research Council, UK (BBSRC): BBFO16980/1, the Medical Research Council, UK (MRC): MR/L001284/1 and the Wellcome Trust: 075294, 093088. The authors would like to thank the editor and anonymous reviewers for their valuable comments and suggestions.

References

Addinall, S. G., M. Downey, M. Yu, M. K. Zubko, J. Dewar, A. Leake, J. Hallinan, O. Shaw, K. James, D. J. Wilkinson, A. Wipat, D. Durocher, and D. Lydall (2008, Dec). A genomewide suppressor and enhancer analysis of *cdc13-1* reveals varied cellular processes influencing telomere capping in *Saccharomyces cerevisiae*. *Genetics* 180(4), 2251–2266.

Addinall, S. G., E.-M. Holstein, C. Lawless, M. Yu, K. Chapman, A. P. Banks, H.-P. Ngo, L. Maringele, M. Taschuk, A. Young, A. Ciesiolka, A. L. Lister, A. Wipat, D. J. Wilkinson, and D. Lydall (2011). Quantitative Fitness Analysis Shows That NMD Proteins and Many Other Protein Complexes Suppress or Enhance Distinct Telomere Cap Defects. *PLoS Genet* 7(4), e1001362.

Aylor, D. L. and Z.-B. Zeng (2008). From classical genetics to quantitative genetics to systems biology: Modeling epistasis. *PLoS Genet* 4(3), e1000029.

Banks, A., C. Lawless, and D. Lydall (2012). A Quantitative Fitness Analysis Workflow. *J. Vis. Exp* 66, e4018.

Benjamini, Y. and Y. Hochberg (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society. Series B (Methodological)* 57(1), 289–300.

Bernardo, J. and A. Smith (2007). *Bayesian Theory*. Wiley Series in Probability and Statistics. John Wiley & Sons Canada, Limited.

Cordell, H. J. (2002). Epistasis: what it means, what it doesn't mean, and statistical methods to detect it in humans. *Human molecular genetics* 11(20), 2463–2468.

Falcon, S. and R. Gentleman (2007). Using GOstats to test gene lists for GO term association. *Bioinformatics* 23(2), 257–8.

Foster, S. S., M. K. Zubko, S. Guillard, and D. Lydall (2006). MRX protects telomeric DNA at uncapped telomeres of budding yeast *cdc13-1* mutants. *DNA Repair* 5(7), 840 – 851.

Franke, J., J. Gehlen, and A. E. Ehrenhofer-Murray (2008, Nov). Hypermethylation of yeast telomerase RNA by the snRNA and snoRNA methyltransferase Tgs1. *J. Cell. Sci.* 121(Pt 21), 3553–3560.

Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models. *Bayesian analysis* 1(3), 515–533.

Gelman, A. and J. Hill (2006). *Data Analysis Using Regression and Multilevel/Hierarchical Models* (1 ed.). Cambridge University Press.

Goldstein, H. (2011). *Multilevel Statistical Models*. Wiley Series in Probability and Statistics. Wiley.

Greider, C. W. and E. H. Blackburn (1985). Identification of a specific telomere terminal transferase activity in tetrahymena extracts. *cell* 43(2), 405–413.

Heidelberger, P. and P. D. Welch (1981). A spectral method for confidence interval generation and run length control in simulations. *Commun. ACM* 24(4), 233–245.

Jow, H., R. J. Boys, and D. J. Wilkinson (2014, July). Bayesian identification of protein differential expression in multi-group isobaric labelled mass spectrometry data. *ArXiv e-prints*.

Keogh, M. C., S. K. Kurdistani, S. A. Morris, S. H. Ahn, V. Podolny, S. R. Collins, M. Schuldiner, K. Chin, T. Punna, N. J. Thompson, C. Boone, A. Emili, J. S. Weissman, T. R. Hughes, B. D. Strahl, M. Grunstein, J. F. Greenblatt, S. Buratowski, and N. J. Krogan (2005, Nov). Cotranscriptional set2 methylation of histone H3 lysine 36 recruits a repressive Rpd3 complex. *Cell* 123(4), 593–605.

Lawless, C., D. J. Wilkinson, A. Young, S. G. Addinall, and D. A. Lydall (2010). Colonyzer: automated quantification of micro-organism growth characteristics on solid agar. *BMC Bioinformatics* 11, 287.

Mani, R., R. P. St.Onge, J. L. Hartman, G. Giaever, and F. P. Roth (2008). Defining genetic interaction. *Proceedings of the National Academy of Sciences* 105(9), 3461–3466.

Nugent, C. I., T. R. Hughes, N. F. Lue, and V. Lundblad (1996). Cdc13p: A Single-Strand Telomeric DNA-Binding Protein with a Dual Role in Yeast Telomere Maintenance. *Science* 274(5285), 249–252.

O'Hara, R. B. and M. J. Sillanpaa (2009). A Review of Bayesian Variable Selection Methods: What, How and Which. *Bayesian Analysis* 4, 85.

Olovnikov, A. (1973, sep). A theory of marginotomy. *Journal of Theoretical Biology* 41(1), 181–190.

Phenix, H., K. Morin, C. Batenchuk, J. Parker, V. Abedi, L. Yang, L. Tepliakova, T. J. Perkins, and M. Krn (2011). Quantitative Epistasis Analysis and Pathway Inference from Genetic Interaction Data. *PLoS Comput Biol* 7(5), e1002048.

Phillips, P. C. (1998). The language of gene interaction. *Genetics* 149(3), 1167–1171.

Pinheiro, J. C. and D. M. Bates (2000). *Mixed Effects Models in S and S-Plus*. Springer.

Raftery, A. E. and S. M. Lewis (1995). The Number of Iterations, Convergence Diagnostics and Generic Metropolis Algorithms. In *In Practical Markov Chain Monte Carlo* (W.R. Gilks, D.J. Spiegelhalter and S. Richardson, eds.), pp. 115–130. Chapman and Hall.

Schuldiner, M., S. Collins, J. Weissman, and N. Krogan (2006). Quantitative genetic analysis in *Saccharomyces cerevisiae* using epistatic miniaarray profiles (E-MAPs) and its application to chromatin functions. *Methods* 40(4), 344 – 352.

Tong, A. H. and C. Boone (2006). Synthetic genetic array analysis in *Saccharomyces cerevisiae*. *Methods Mol Biol* 313, 171–192.

Verhulst, P. F. (1845). Recherches mathématiques sur la loi d'accroissement de la population. *Nouveaux mémoires de l'Academie Royale des Science et Belles-Lettres de Bruxelles* 18, 1–41.

Yi, N. (2010). Statistical analysis of genetic interactions. *Genetics research* 92(5-6), 443–459.

Zhang, L., V. Baladandayuthapani, B. K. Mallick, G. C. Manyam, P. A. Thompson, M. L. Bondy, and K.-A. Do (2014). Bayesian hierarchical structured variable selection methods with application to molecular inversion probe studies in breast cancer. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 63(4), 595–620.

Web-based supporting materials for “Bayesian hierarchical modelling for inferring genetic interactions in yeast” by Jonathan Heydari, Conor Lawless, David A. Lydall and Darren J. Wilkinson

1 Plate diagrams

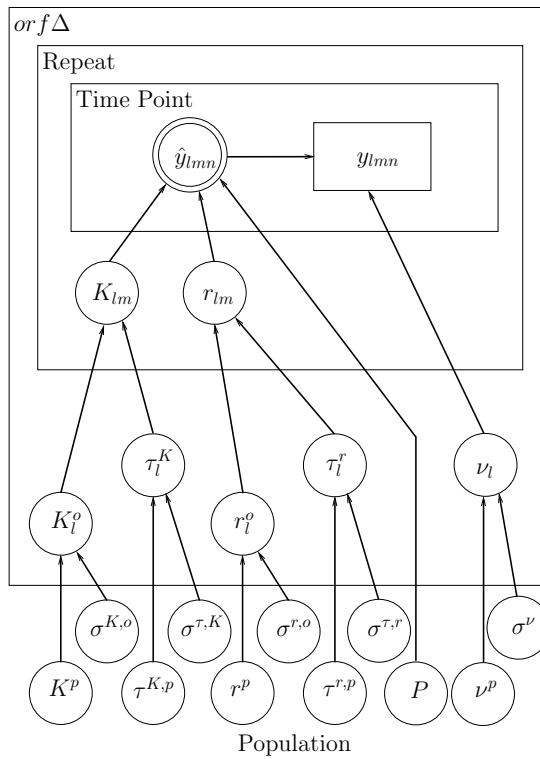


Figure 1.1: Plate diagram for the SHM, described in Section 3.1 of the main article. This figure shows the four levels of hierarchy in the SHM model, population, $orf\Delta$ (l), repeat (m) and time point (n). Prior hyperparameters for the population parameters are omitted. A circular node represents a parameter in the model. An arrow from a source node to a target node indicates that the source node parameter is a prior hyperparameter for the target node parameter. Each rectangular box corresponds to a level of the hierarchy. Nodes within multiple boxes are nested and their parameters are indexed by corresponding levels of the hierarchy. The node consisting of two concentric circles corresponds to our model’s fitted values. The rectangular node represents the observed data.

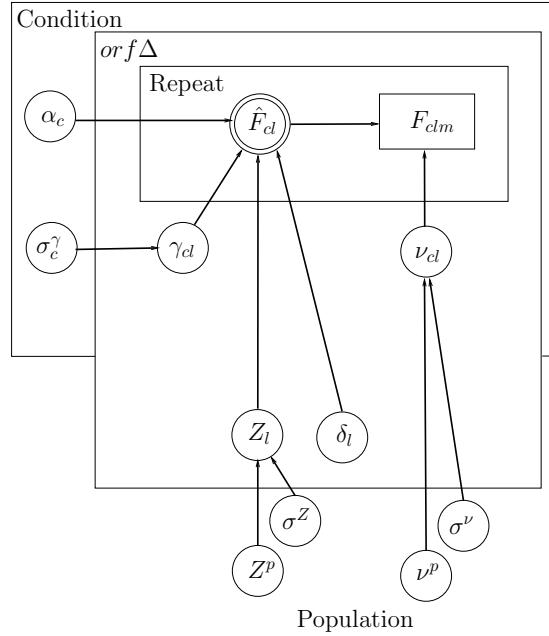


Figure 1.2: Plate diagram for the IHM, described in Section 3.2 of the main article. This figure shows the four levels of hierarchy in the IHM model: population, $orf\Delta$ (l), condition (c) and repeat (m). Prior hyperparameters for population parameters are omitted. Plate diagram notation as in Figure 1.1.

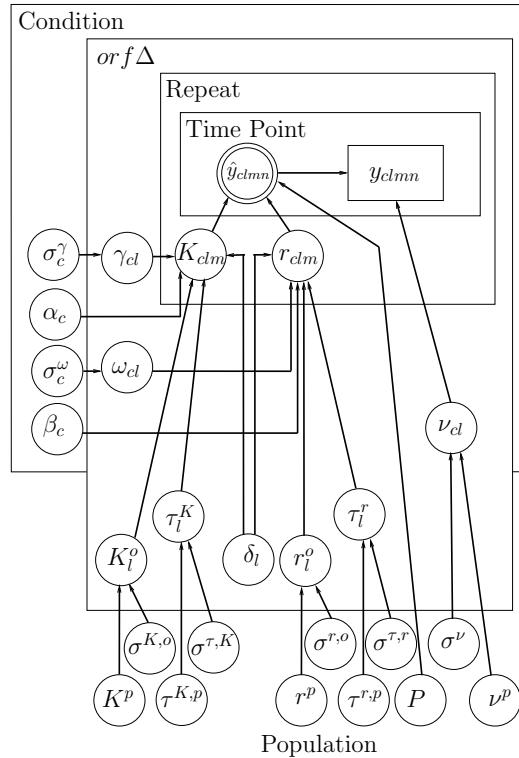


Figure 1.3: Plate diagram the JHM, described in Section 3.3 of the main article. This figure shows the five levels of hierarchy in the JHM model, population, $orf\Delta$ (l), condition (c), repeat (m) and time point (n). Prior hyperparameters for the population parameters are omitted. Plate diagram notation is given in Figure 1.1.

2 GO term enrichment analysis in R

```
source("http://bioconductor.org/biocLite.R")
biocLite("GOSTats")
biocLite("org.Sc.sgd.db")
#####
library(GOSTats) # GO testing tool package
library(org.Sc.sgd.db) # yeast gene annotation package
genes=read.table("sm_JHM_list.txt", header=T)
UNIVSTRIP=genes[,2]
genes<-as.vector(genes[genes[,3]>0.5,2])
genes<-unique(genes)
ensemblIDs=as.list(org.Sc.sgdPMID2ORF)
univ=unlist(ensemblIDs)
univ=univ[!is.na(univ)]
length(univ)
length(unique(univ))
univ=unique(univ)
all=as.vector(univ)
all=all[all%in%UNIVSTRIP]
length(all)
ontology=c("BP")
vec<-genes%in%univ
genes<-genes[vec]
params_temp=new("GOHyperGParams", geneIds=genes,
universeGeneIds=all,
annotation="org.Sc.sgd.db", categoryName="GO",
ontology=ontology, pvalueCutoff=1,
testDirection = "over")
results=hyperGTest(params_temp)
results=summary(results)
results$qvalue<-p.adjust(results$pvalue,method="BH")
```

3 Fitness plots with GO terms highlighted

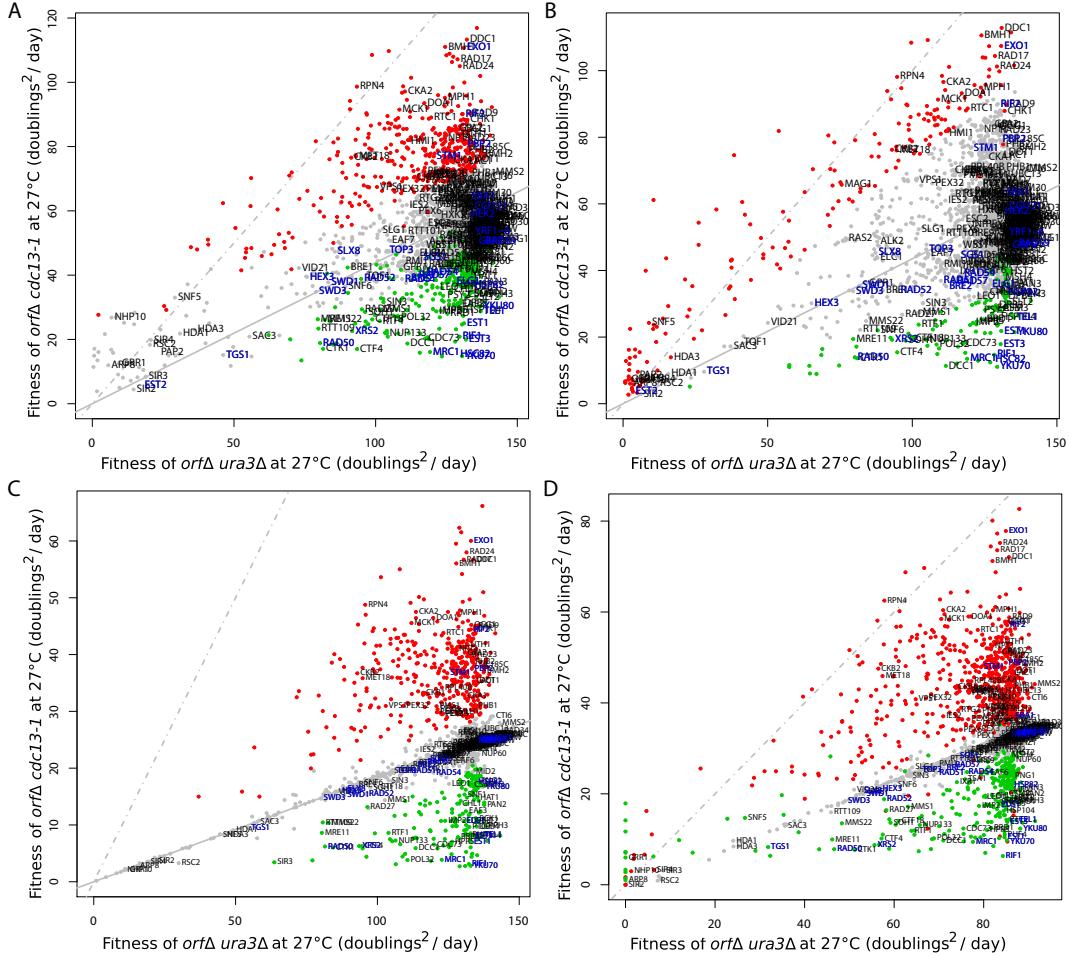


Figure 3.1: Alternative fitness plots with $orf\Delta$ posterior mean fitnesses. Text for the “telomere maintenance” GO term is highlighted in blue. A) Non-Bayesian, non-hierarchical fitness plot, based on Table S6 from Addinall et al. (2011) ($F = MDR \times MDP$). B) Non-Bayesian, hierarchical fitness plot, from fitting REM to data in Table S6 from Addinall et al. (2011) ($F = MDR \times MDP$). C) IHM fitness plot with $orf\Delta$ posterior mean fitness ($F = MDR \times MDP$). D) JHM fitness plot with $orf\Delta$ posterior mean fitnesses. $orf\Delta$ strains are classified as being a suppressor or enhancer based on analysis of growth parameter r . Further explanation and notation for fitness plots are given in Figure 3 of the main article.

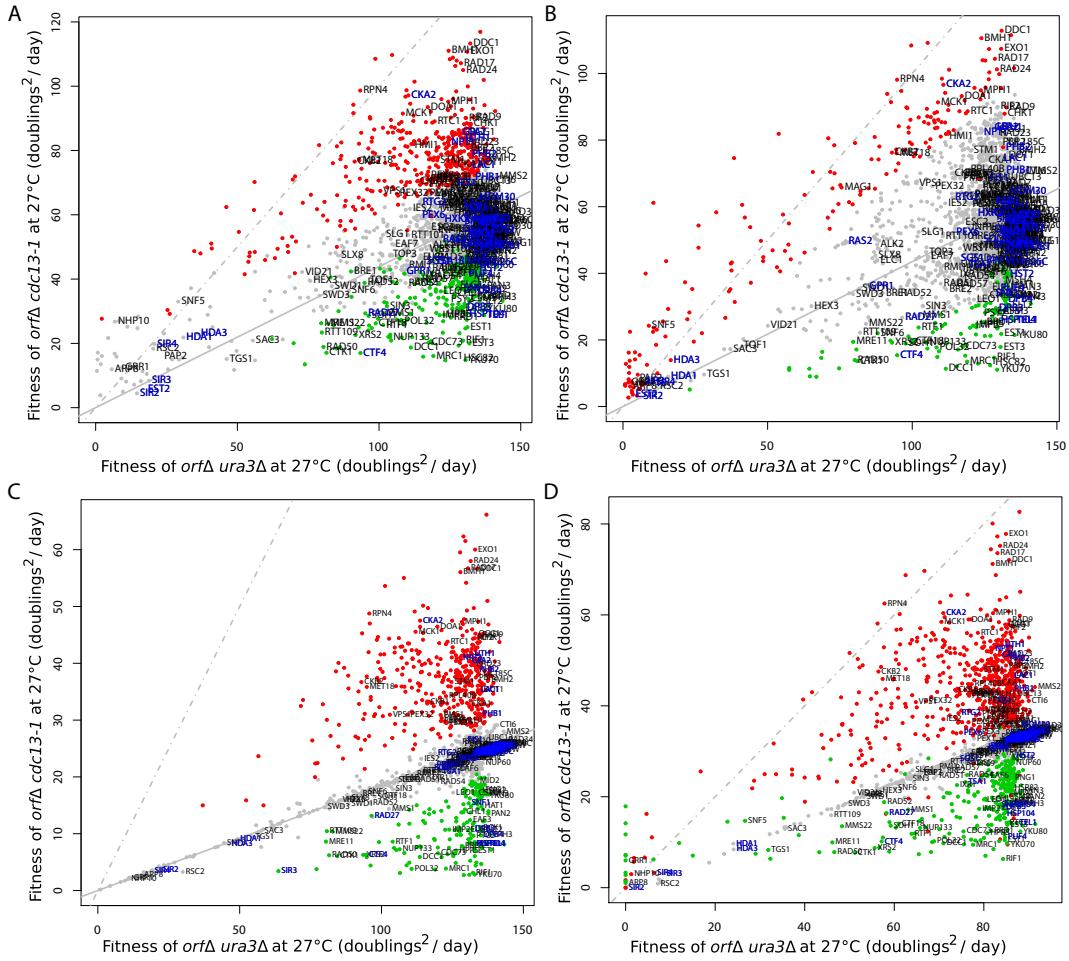


Figure 3.2: Alternative fitness plots with $orf\Delta$ posterior mean fitnesses. Text for the “ageing” GO term is highlighted in blue. A) Non-Bayesian, non-hierarchical fitness plot, based on Table S6 from Addinall et al. (2011) ($F = MDR \times MDP$). B) Non-Bayesian, hierarchical fitness plot, from fitting REM to data in Table S6 from Addinall et al. (2011) ($F = MDR \times MDP$). C) IHM fitness plot with $orf\Delta$ posterior mean fitness ($F = MDR \times MDP$). D) JHM fitness plot with $orf\Delta$ posterior mean fitnesses. $orf\Delta$ strains are classified as being a suppressor or enhancer based on analysis of growth parameter r . Further explanation and notation for fitness plots are given in Figure 3 of the main article.

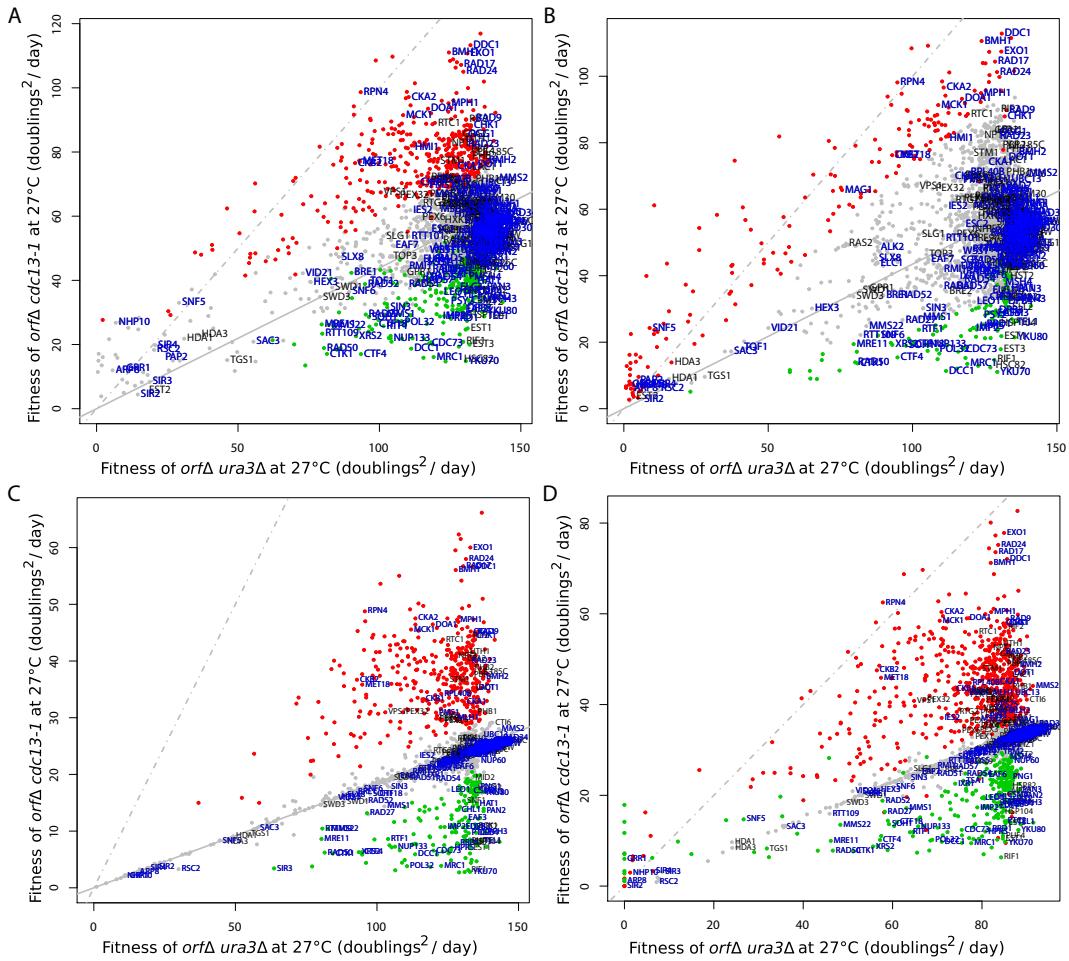


Figure 3.3: Alternative fitness plots with $orf\Delta$ posterior mean fitnesses. Text for the “response to DNA damage” GO term is highlighted in blue. A) Non-Bayesian, non-hierarchical fitness plot, based on Table S6 from Addinall et al. (2011) ($F = MDR \times MDP$). B) Non-Bayesian, hierarchical fitness plot, from fitting REM to data in Table S6 from Addinall et al. (2011) ($F = MDR \times MDP$). C) IHM fitness plot with $orf\Delta$ posterior mean fitness ($F = MDR \times MDP$). D) JHM fitness plot with $orf\Delta$ posterior mean fitnesses. $orf\Delta$ strains are classified as being a suppressor or enhancer based on analysis of growth parameter r . Further explanation and notation for fitness plots are given in Figure 3 of the main article.

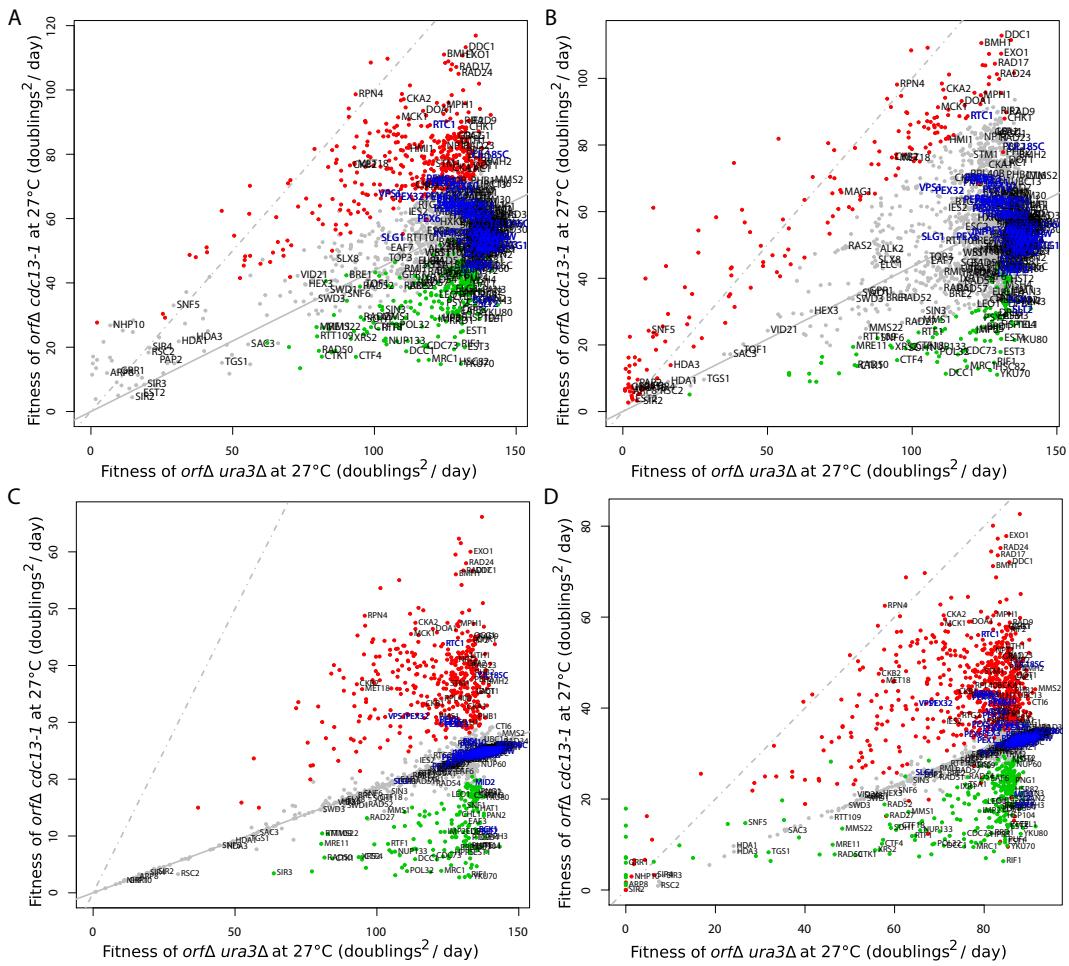


Figure 3.4: Alternative fitness plots with $orf\Delta$ posterior mean fitnesses. Text for the “peroxisomal organisation” GO term is highlighted in blue. A) Non-Bayesian, non-hierarchical fitness plot, based on Table S6 from Addinall et al. (2011) ($F = MDR \times MDP$). B) Non-Bayesian, hierarchical fitness plot, from fitting REM to data in Table S6 from Addinall et al. (2011) ($F = MDR \times MDP$). C) IHM fitness plot with $orf\Delta$ posterior mean fitness ($F = MDR \times MDP$). D) JHM fitness plot with $orf\Delta$ posterior mean fitnesses. $orf\Delta$ strains are classified as being a suppressor or enhancer based on analysis of growth parameter r . Further explanation and notation for fitness plots are given in Figure 3 of the main article.

4 Lists of top genetic interactions for IHM and JHM approaches

<i>Type of Interaction</i>	<i>Gene Name</i>	<i>Probability of Interaction</i> δ_l	<i>Strength of Interaction</i> $e^{(\delta_l \gamma_l)}$	<i>Position in Addinall (2011)</i>
Suppressor	IPK1	1.00	2.87	10
	LST4	1.00	2.77	13
	RPN4	1.00	2.76	17
	MTC5	1.00	2.66	20
	GTR1	1.00	2.64	38
	NMD2	1.00	2.62	3
	SAN1	1.00	2.62	16
	UPF3	1.00	2.58	21
	RPL37A	1.00	2.56	121
	NAM7	1.00	2.53	22
	RPP2B	1.00	2.52	120
	YNL226W	0.99	2.49	126
	YGL218W	1.00	2.46	250
	MEH1	1.00	2.45	45
	ARO2	1.00	2.45	68
	EXO1	1.00	2.45	1
	BUD27	1.00	2.43	46
	RAD24	1.00	2.39	4
	RPL16B	1.00	2.39	33
	RPL43A	1.00	2.39	150
Enhancer	MRC1	1.00	0.11	35
	YKU70	1.00	0.11	31
	STI1	1.00	0.11	42
	RIF1	1.00	0.13	36
	ELP3	1.00	0.16	82
	CLB5	1.00	0.17	58
	MRC1	1.00	0.17	63
	DPH2	1.00	0.18	24
	POL32	1.00	0.19	113
	MAK31	1.00	0.19	37
	SWM1	1.00	0.20	25
	LTE1	1.00	0.21	48
	MAK10	1.00	0.22	44
	ELP2	1.00	0.22	77
	PAT1	1.00	0.24	144
	DPH1	1.00	0.25	55
	SRB2	0.99	0.25	174
	THP2	1.00	0.26	67
	MFT1	1.00	0.26	52
	LSM6	0.97	0.26	389

A file containing the full list of genetic interactions is also provided in the on-line supporting materials.

Table 4.1: Sample of IHM top genetic interactions

Type of Interaction	Gene Name	Probability of Interaction δ_l	Strength of Interaction $e^{(\delta_l \gamma_l)}$	Strength of Interaction $e^{(\delta_l \omega_l)}$	Strength of Interaction $MDR \times MDP$	Position in Addinall (2011)
Suppressor in K	CSE2	1.00	490.51	0.48	11.71	838
	SGF29	1.00	273.69	0.68	14.16	580
	GSH1	1.00	78.79	0.92	17.89	281
	YMD8	1.00	59.31	0.65	7.05	2022
	YGL024W	1.00	28.13	1.18	13.33	151
	RPS9B	1.00	24.67	1.12	10.24	801
	GRR1	1.00	22.51	0.67	5.99	1992
Suppressor in r	BTS1	1.00	19.27	2.29	19.65	201
	IPK1	1.00	5.56	2.26	44.81	10
	NMD2	1.00	2.96	2.19	48.51	3
	SAN1	1.00	2.37	2.17	48.70	16
	LST4	1.00	5.79	2.14	44.14	13
	RPN4	1.00	8.00	2.12	40.46	17
	UPF3	1.00	3.16	2.07	45.25	21
Suppressor in $MDR \times MDP$	SAN1	1.00	2.37	2.17	48.70	16
	NMD2	1.00	2.96	2.19	48.51	3
	UPF3	1.00	3.16	2.07	45.25	21
	EXO1	1.00	2.89	2.06	45.04	1
	IPK1	1.00	5.56	2.26	44.81	10
	LST4	1.00	5.79	2.14	44.14	13
	NAM7	1.00	3.02	2.04	43.00	22
Enhancer in K	YKU70	1.00	0.01	1.09	-23.44	31
	STI1	1.00	0.01	1.20	-21.60	42
	RIF1	1.00	0.01	0.63	-26.17	36
	MRC1	1.00	0.01	0.83	-23.15	35
	MAK31	1.00	0.02	1.18	-18.19	37
	CLB5	1.00	0.02	0.87	-19.54	58
	MRC1	1.00	0.02	0.81	-20.40	63
Enhancer in r	PAT1	1.00	1.71	0.28	-18.30	144
	PUF4	1.00	2.00	0.31	-21.61	34
	YKU80	1.00	2.15	0.33	-21.68	32
	RTT103	1.00	2.54	0.34	-17.87	153
	LSM1	0.99	2.13	0.34	-16.20	101
	GIM3	0.99	0.93	0.35	-19.70	132
	INP52	0.96	0.86	0.36	-14.50	345
Enhancer in $MDR \times MDP$	RIF1	1.00	0.01	0.63	-26.17	36
	LTE1	1.00	0.06	0.40	-23.96	48
	YKU70	1.00	0.01	1.09	-23.44	31
	MRC1	1.00	0.01	0.83	-23.15	35
	DPH2	1.00	0.04	0.56	-23.11	24
	EST1	1.00	0.12	0.46	-22.20	5
	MAK10	1.00	0.04	0.59	-21.92	44

A file containing the full list of genetic interactions is also provided in the on-line supporting materials.

Table 4.2: Sample of JHM top genetic interactions

5 Alternative fitness plots for the JHM

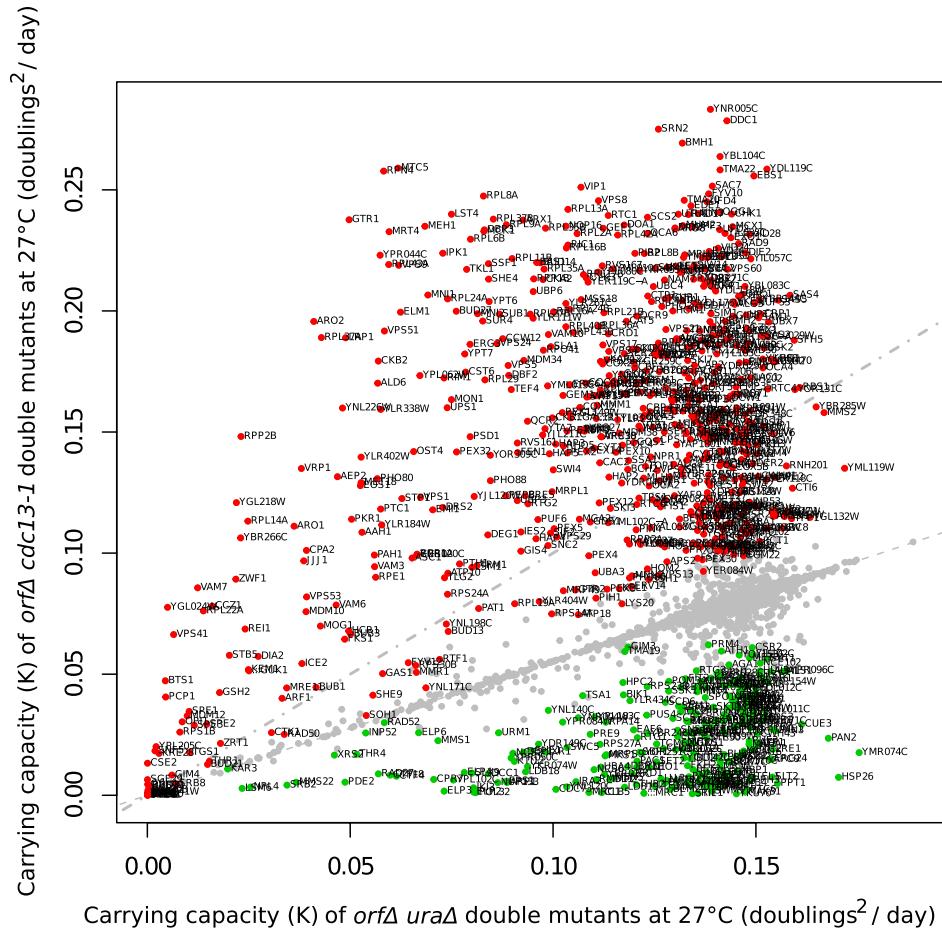


Figure 5.1: Joint hierarchical model (JHM) carrying capacity fitness plot with $orf\Delta$ posterior mean fitnesses. $orf\Delta$ strains are classified as being a suppressor or enhancer based on carrying capacity parameter K . Further explanation and notation for fitness plots are given in Figure 3 of the main article.

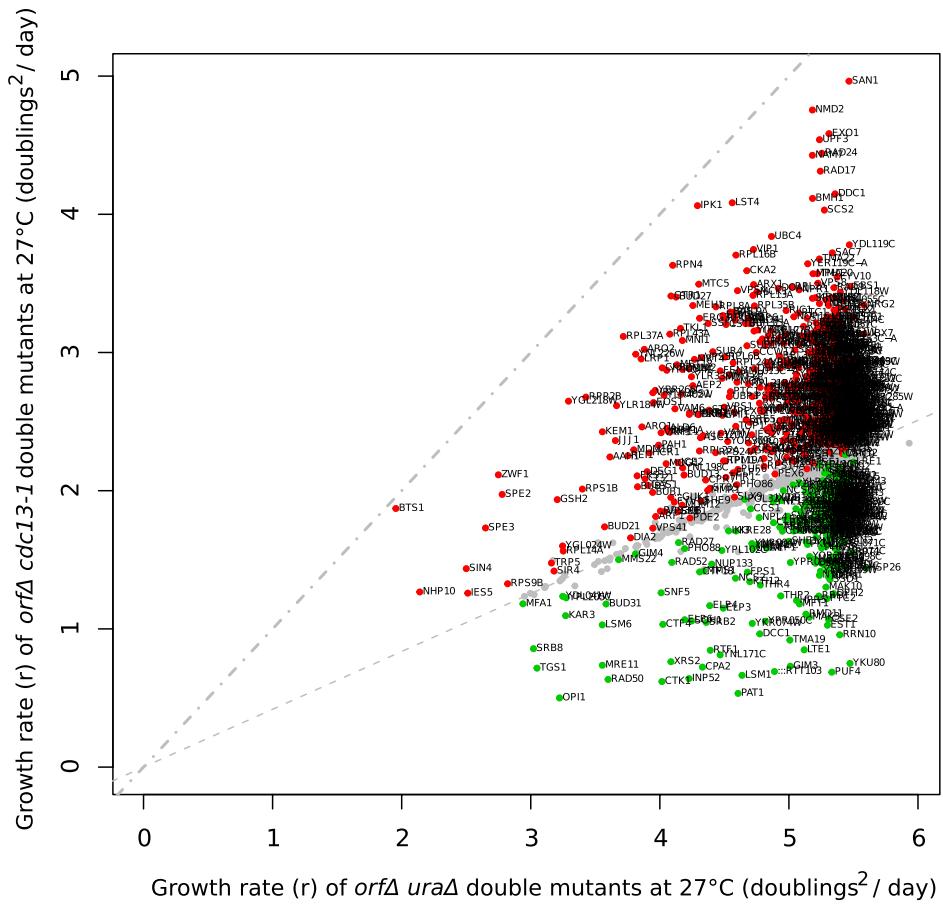


Figure 5.2: Joint hierarchical model (JHM) growth rate fitness plot with $orf\Delta$ posterior mean fitnesses. $orf\Delta$ strains are classified as being a suppressor or enhancer based on growth parameter r . Further explanation and notation for fitness plots are given in Figure 3 of the main article.

6 The effect of parameter p on specificity

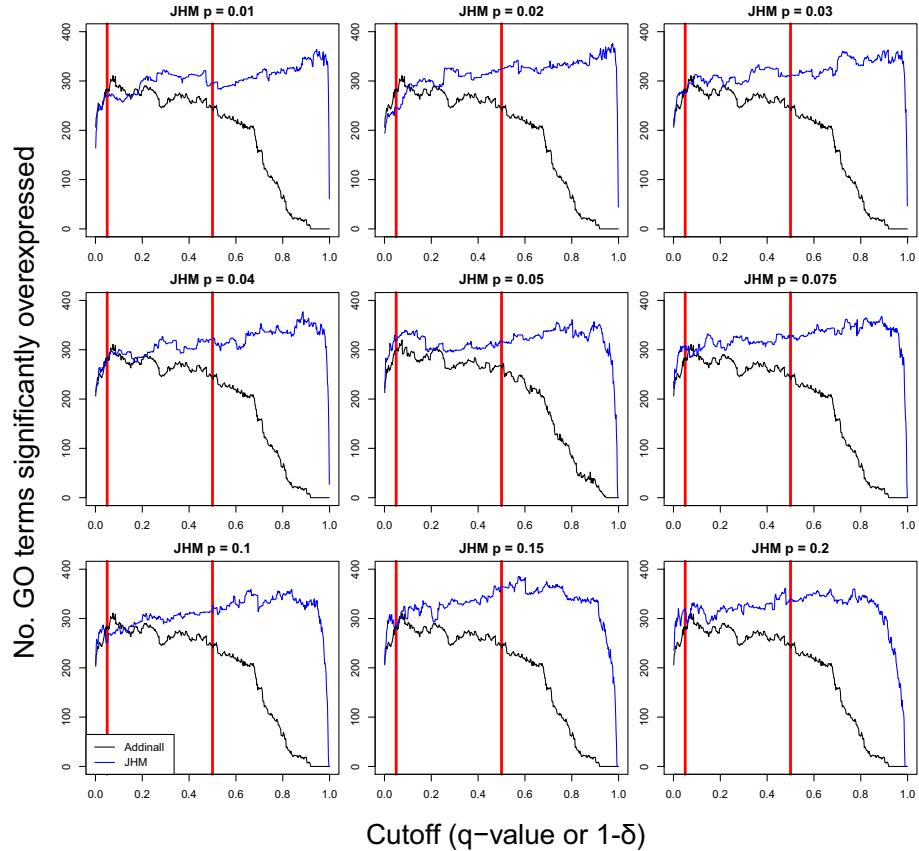


Figure 6.1: Comparison of the number of significantly over-expressed GO terms identified in lists of significant interactors found using the Addinall (2011) method and using the JHM. Significantly over-expressed GO terms were identified using the hyperGTest function in the GOstats R package. Note that the values used to classify whether a gene interacts with *cdc13-1* at 27C (q-value and δ respectively, red vertical lines as presented in Section 4.4) are not directly comparable. However, the full range of possible cutoffs for both values are plotted. Each panel shows the change in over-expressed GO terms with cutoff for a different value of the p parameter (prior estimate of expected proportion of interactors) used in the JHM analysis.

7 Correlation between methods

The Addinall et al. (2011) approach has its highest correlation with the IHM, followed by the JHM and then the REM. The REM correlates least well with the JHM while showing the same correlation with both the Addinall et al. (2011) approach and the IHM. The correlation between the IHM and the JHM is the largest observed between any of the methods, demonstrating the similarity of our Bayesian hierarchical methods.

<i>Method</i>	<i>Method</i>			
	<i>Addinall et al. (2011)</i> <i>QFA</i>	<i>REM</i> <i>QFA</i>	<i>IHM</i> <i>QFA</i>	<i>JHM QFA</i> <i>(MDR × MDP)</i>
Addinall et al. (2011) QFA,	1	0.77	0.89	0.88
REM QFA,		1	0.77	0.75
IHM QFA,			1	0.95
JHM QFA (<i>MDR × MDP</i>),				1

Table 7.1: Spearman’s rank correlation coefficients for magnitudes from genetic independence, between Addinall et al. (2011), REM, IHM and JHM QFA methods

The $MDR \times MDP$ correlation plot of the JHM versus the Addinall et al. (2011) approach demonstrates the similarity (Pearson correlation=0.90) and differences between the two approaches in terms of $MDR \times MDP$. We can see how the results differ between the JHM and Addinall et al. (2011), with a kink at the origin due to the JHM allowing shrinkage of non-interacting genes towards the fitted line.

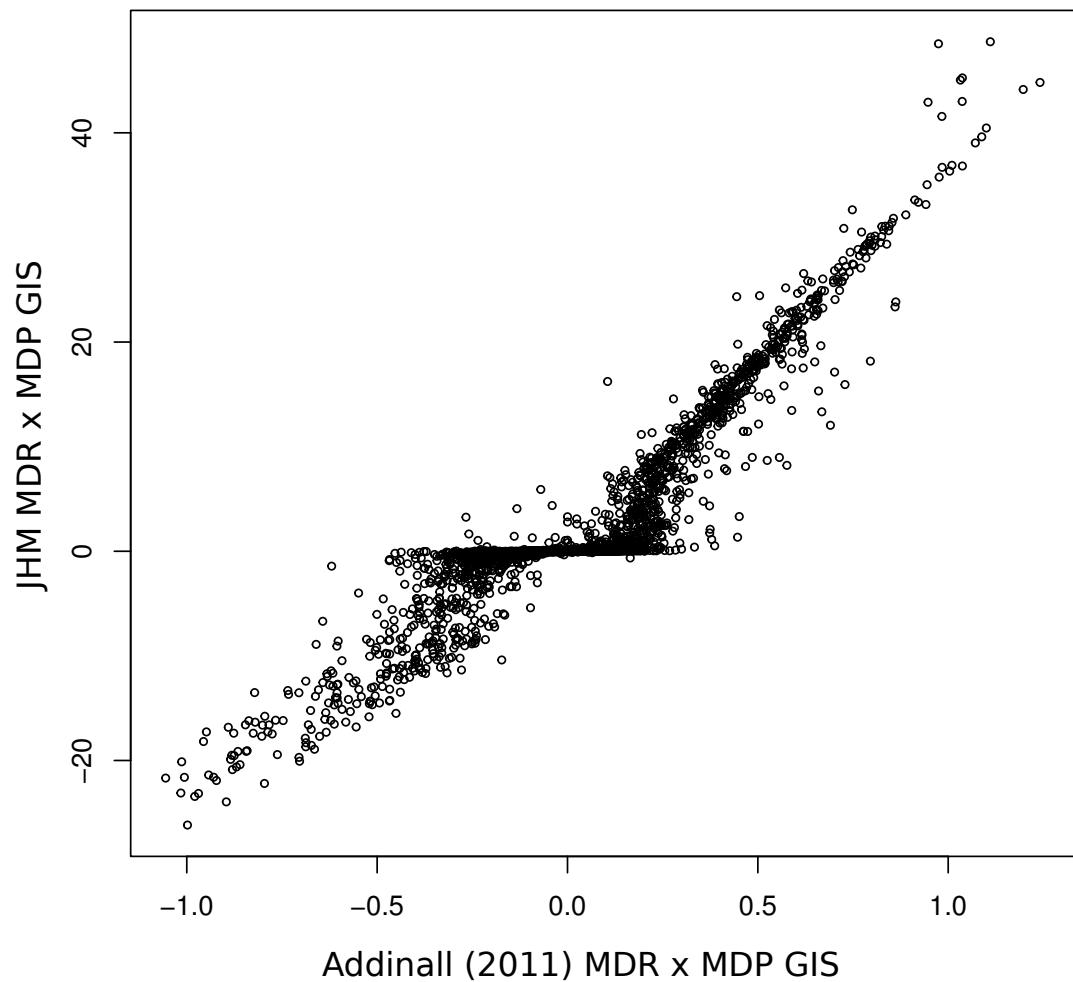


Figure 7.1: $MDR \times MDP$ genetic interaction correlation plot of JHM versus Addinall et al. (2011) (Pearson correlation=0.90).

8 A simulation study comparing specificity and sensitivity of the Addinall et al. (2011) approach, the SHM and the JHM

Since our understanding of biological processes is currently incomplete it is difficult to assess what proportion of genetic interactions identified by fitting any model to real biological data are real. If we don't know which interactions are true, we cannot know which of those interactions identified by any inference scheme are false positives. In order to compare the ability of each of our models to identify subtle, true interactions (sensitivity) while avoiding false positives (specificity), a separate simulation study was carried out (Section 4.3.6 <http://arxiv.org/abs/1405.7091>). Synthetic control and query datasets of similar size, quality and resolution to real QFA datasets, with known suppressors and enhancers of a simulated query mutation were constructed using a hierarchical simulation model consistent with the JHM. We used the JHM to simulate the synthetic dataset since it is the most detailed model we have available and the one which most closely matches the structure of QFA experiments. The Addinall et al. (2011) approach, the REM, the SHM and the JHM were each fit to the synthetic dataset and the lists of suppressors and enhancers as well as the list of all interactors generated by each method were compared with the list of known true interactors.

Sensitivity and specificity achieved with each of the models were presented in table 4.4 <http://arxiv.org/abs/1405.7091>. In summary, the simulation study showed that the JHM correctly identified a higher proportion of true interactions (314/430) than the Addinall et al. (2011) approach (220/430), while also identifying fewer false positives (JHM: 8, Addinall et al. (2011): 303).